Journal of Directed Energy, **4**, Winter 2010, 168–188

Characterization of High-Power Lasers

Jack Slater*

Schafer Corporation, 2309 Renard Place, SE, Suite 300, Albuquerque, New Mexico 87106

A general methodology for characterization of high-average-power lasers with respect to power and beam quality is presented. The techniques discussed largely capture the experience from the 100 kW Joint High Power Solid State Laser (JHPSSL) program and plans for the Robust Electric Laser Initiative (RELI) program.

KEYWORDS: Beam quality, High-power laser, Laser, Laser characterization

1. Overview

This technical note discusses measurement methods for characterization of high-power lasers. It is intended to promote a common measurement approach for high-energy laser (HEL) characterization in the directed energy (DE) Department of Defense (DOD) community. Such uniformity will improve the efficiency of dialogue between laser contractors, those evaluating various DE applications, and Government sponsors.

The methodology and technical details provided here result largely from experience in the 25- and 100-kW phases of the HEL-JTO JHPSSL (High Energy Laser Joint Technology Office–Joint High Power Solid State Laser) program and planning for the current RELI (Robust Electric Laser Initiative) program. Development of technical requirements for these programs has involved interaction with numerous contractors and Government research groups in order to provide a workable set of metrics. This document supersedes a similar characterization note that was written at the start of the JHPSSL program in 2004,[3] and the present iteration coincides with the start of the RELI program.

The paper is written at a fairly abstract level and in a somewhat tutorial fashion, concentrating on the underlying measurement principles, explanation of what we consider to be good practice, and identification of areas that need special attention. It is intended for those already familiar with laser characterization. More general information regarding HEL characterization can be found in Refs. 5 and 6 and references therein. One purpose of this paper is to promote dialogue between laser contractor teams and those who will perform independent measurements at a program milestone or conclusion. It was borne out in the JHPSSL program that if both groups adopt common language and measurement methods, then the contractors can develop characterization systems that will produce measurement values consistent with those of an independent measurement team using different hardware. Reliable and reproducible results provide a solid backdrop for top-level programmatic

Received December 6, 2010

*Corresponding author; e-mail:jslater@schaferalb.com.

decisions. A second purpose of this paper is to provide a background for writing detailed laser specifications.

The measurements described here do not provide the level of detail that may be needed for, as an example, preparing an interface specification between a laser and a beam director. For that, a near-field intensity distribution would likely be needed and possibly also a map of the photon angular distribution such as might be measured by a wavefront sensor.

In contrast, the approach presented here is intended to utilize the simplest equipment and the simplest calibration in order to provide a useful measurement that has the least chance of unanticipated complications. Specifically, this approach uses physical pinholes (which may be augmented by cameras), i.e., a hard-aperture approach, and does not require measurement of intensity or angular distribution functions. The intent is to provide robust measurements in the sense that two groups would agree on measurement of a given laser beam, e.g., a contractor and an independent measurement team, in order that program management needs can be served.

Captured here is one characterization path that has been successful. Certainly other measurement strategies are possible and may be more desirable when more detailed information is needed.

The organization is as follows. Section 2 describes measurement of laser power based on use of a large, known attenuation of the HEL beam in conjunction with an absolutely calibrated, low-power detector. Section 3 describes power in the bucket (PIB) beam quality (BQ) measurement methods. Section 4 describes measurement of other parameters, such as jitter and start-up time, and Section 5 outlines an example characterization-report template.

2. Measurement of Output Power

A basic measurement approach that has been used quite successfully up to the 100-kW level is shown in Fig. 1. There are essentially three pieces:

(a) a known attenuation, of order 30–40 dB, which captures a low-power sample while passing most of the beam to a dump that is not calibrated.
(b) some intervening optics, with modest loss, e.g., 3 dB, whose loss is calibrated with the aid of a sensor of known linearity.

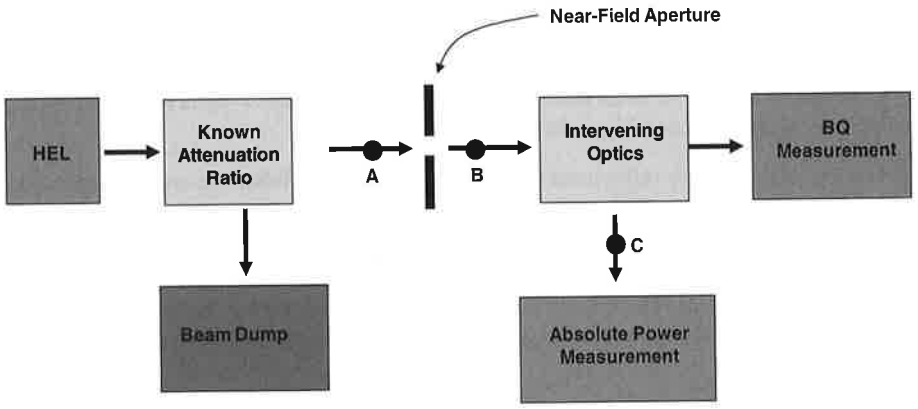

Fig. 1. Power measurement based on use of a known attenuation and a low-power sensor with absolute calibration.

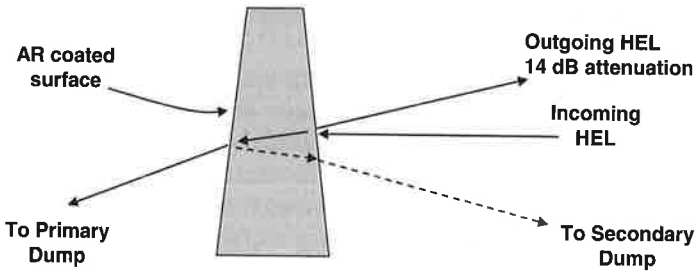

Fig. 2. One stage of front-surface attenuation based on a known Fresnel reflection coefficient. One secondary reflection is shown as a dashed line; others exist.

(c) a sensor that has been absolutely calibrated by NIST (National Institute of Standards and Technology).

This arrangement has the benefit that it is well suited for portable hardware that will be transported to an established laser facility where an existing beam dump can be used.

An alternative to the scheme of Fig. 1 is a measurement made at full power using a calorimeter. Such measurements can be made reasonably accurately with ease, but they tend to suffer from lack of temporal bandwidth and hence may be less useful for HEL pulse trains of interest.

Each of these elements of Fig. 1 has considerations as described below.

2.1. Known attenuation ratio

Good agreement between measurement teams was obtained in the JHPSSL program using multiple stages of attenuation via front-surface reflection from uncoated fused silica. The uncoated surface is used so that the reflection coefficient can be calculated using the known index of refraction, wavelength, polarization, and angle of incidence. Use of an uncoated surface also eliminates the need to assess any coating reflectivity variability due to heating effects at high power. For near-normal incidence, each reflection is roughly a 14-dB loss so that, for example, only four stages are needed to reduce a 100-kW beam to levels at which the beam is easily manipulated, even with neutral density filters. Figure 2 shows a single stage of attenuation. An antireflection (Ar) coating on the rear surface of this wedge is useful to reduce the power contained in secondary beam(s).

In the case in which the calibrated attenuation is a series of front surface reflections, e.g., two or four, there are several key challenges to meet:

(a) Control of secondary reflections to avoid scattered light. When the overall attenuation is large, e.g., of order 60 dB, even small levels of scattering or secondary reflections can interfere with power and beam quality measurements. Careful use of shrouds and baffles is required. For a case with 100-kW incident power and an AR coating reflectivity of 0.1%, an unwanted, 100-W secondary beam results that may compromise analysis of the *signal* beam, which would be of order 250 mW if four stages of 14-dB attenuation were used.

(b) Control of secondary reflections to avoid air turbulence. There are two air turbulence mechanisms to avoid: (1) heating of the air directly by the high-power beam and (2) secondary air heating via heated surfaces. One must consider the possibility that

surfaces may be heated by backscatter from the primary 100-kW beam dump, and the secondary dump (for the 100-W beam) is also a heat source.

(c) Avoidance of thermal distortion in the first wedge (Fig. 2). It is of paramount importance that the first attenuation wedge (i.e., the wedge passing the HEL beam) has low absorption in order that thermally induced distortion is avoided. This wedge is heated by front-surface absorption (that may be worsened by surface contaminants), bulk absorption, and absorption of the AR coating on the rear surface. The AR coating absorption, typically of order 10 ppm, may dominate overall heat load. An example of very-low-absorption fused silica is Hereaus SUPRASIL 3001 with a rated bulk absorption of 0.25 ppm/cm at 1-μm wavelength. This material is suitable for first-wedge use for at least the 100-kW continuous-wave (CW) range based on the JHPSSL experience. It is worth noting that the thermal gradients developed in the wedge are generally the source of distortions; the average temperature is a problem only if the mount cannot accommodate thermal growth. The coupling of absorption to distortion in cases such as this was analyzed in the Ph.D. thesis of Ryan Lawrence.[4]

(d) Control of polarization effects. It is generally desirable that the calibrated attenuation is independent of the polarization state of the HEL. The angle of incidence will be at least a few degrees off normal, and hence reflectivity will differ for S and P polarization. One compensation for this is to use an even number of attenuation wedges, say, two, and use an out-of-plane geometry such that the S-polarized photons on one wedge become the P-polarized photons on the other.

2.2. Absolutely calibrated power meter

Referring again to Fig. 1 and the absolutely calibrated power meter, this meter preferably has direct NIST calibration rather than relying on a vendor's claim of NIST traceability. We are interested in validating power in pulse trains relevant to DE, which we assume here includes pulsed waveforms. Often the calibrated detectors available in HEL laboratories are based on thermal response and hence may be unable to follow pulses. A particular variation on Fig. 1 would be that there are actually two power meters. One is a low-bandwidth detector (e.g., thermopile) that is absolutely calibrated. The other would operate at higher bandwidth (e.g., a photodiode) and would be calibrated using the slower detector as a standard. If the HEL can operate with long pulses at some reduced power level, then it can be used to cross calibrate the faster detector with adjustment of the known attenuation if needed. If the power meter operates at low power because the known attenuation is large, a low-power surrogate laser could be used for cross calibration if any differences between it and the HEL are accounted for, e.g., wavelength.

2.3. Measuring loss of intervening optics

Referring yet again to Fig. 1, we are assuming that there may be a significant number of intervening optics that provide manipulation in order that the large-diameter HEL beam matches the smaller optics of the low-power diagnostics. These optics may include splitters, lenses, polarization filters, and neutral density filters. Certainly the attenuation of such optics must be measured, at least for those in the path to the power meter. If a power meter of known linearity is available, these losses can be measured directly by using this single meter at positions B and then C while a probe beam is inserted at position A. When using

a probe, one must take into account how probe properties, e.g., wavelength, polarization state, and collimation, might vary from the HEL under testing.

3. Measurement of Beam Quality

3.1. General comments

We begin by noting that a laser beam occupies an $A\Omega$ *product* space defined by its *near-field area*, A (taken to be at a plane of collimation), times its *far-field solid angle*, Ω. This product is conserved in an unaberrated optical system. The units could be mm^2 $mrad^2$. But most often only one plane is considered at a time, and in that case the product space could be measured in $mm\ mrad$.

There is no standard definition of the term *beam quality*. One usage is that BQ is a dimensionless number, of unity or greater, which is the ratio of the beam's $A\Omega$ product compared to that of an ideal beam that is diffraction limited (DL). As above, this usually refers to only one plane at a time. So, for example, BQ = 2, or 2 × DL, would refer to a beam that, from a given beam director (i.e., a given diameter), spreads at twice the angle of the DL beam. Equivalently, at the target it has twice the diameter of the DL beam. Below, we will explain that this is a "horizontal" concept and that there is also a "vertical" concept that indicates what fraction of the beam power would fall into some specified spot size on the target.

Beams in general do not have well-defined edges, so that any measure of the occupied $A\Omega$ product requires a definition of what constitutes the size of the beam, in both physical space and angle space. For the simple and well-defined case of a DL, untruncated Gaussian, and when the size in physical space (i.e., near field) and angle space (i.e., far field) are both measured by the e^{-2} points in intensity, then the AΩ space occupied in one plane, $\pi r\theta$, is exactly λ. For a 1-μm wavelength, this gives 1 mm mrad.

Often the difference from one BQ definition to another stems from differences in how the size of the beam is defined. It could be the second-moment value (e.g., as in Ref. 6 in connection with the M-squared BQ definition) or the full width at half-maximum, or the e^{-2} point, etc. We note that the second-moment definition of size has sensitivity to the far edges of the distribution function, making it poorly suited to DE applications. The PIB approach to beam sizes and distribution functions is described in the next subsection.

3.2. Power-in-the bucket concepts

At this time the DE community is comfortable with PIB program requirements, and measurement methods related to it are described here. PIB methods are well suited to our goal of simple, reproducible measurements. First, some terminology:

PIB curve: Shows the distribution of power in the far field. If azimuthal symmetry is assumed, then this curve is the fractional power enclosed as a function of solid angle.

Vertical BQ measurement method: Is a single-parameter description of *goodness* in which a far-field spot size is specified and the fractional power falling within it is measured and compared to that for an ideal beam. This metric is most meaningful when the beam is nearly DL.

Horizontal BQ measurement method: Is a single-parameter description of goodness in which some fractional power is specified (e.g., 75%) and the far-field spot size containing that power is measured and compared to the spot size for an ideal beam. This metric can be useful even when the beam is much worse than DL.

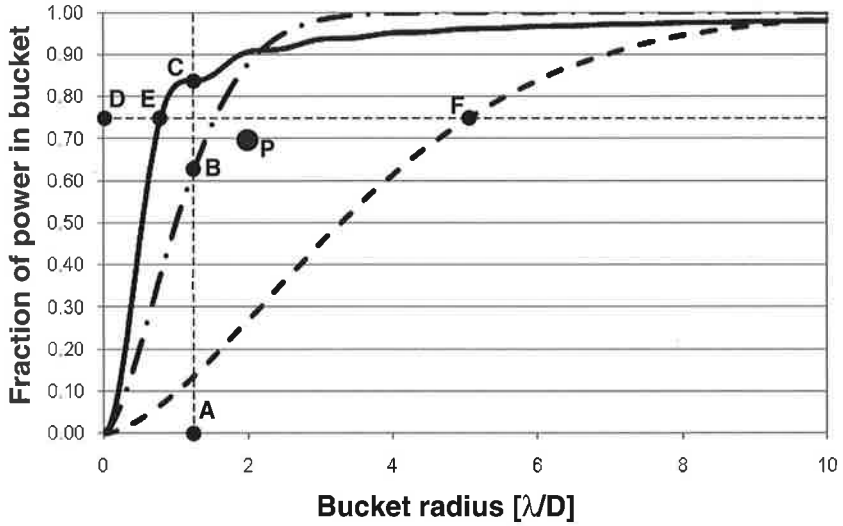

Fig. 3. Example PIB curves. The solid line is a diffraction limited *ideal* PIB curve calculated for a uniformly illuminated disk with uniform phase. In this case the far field pattern is an Airy disk (not shown), and the size of the central diffraction lobe is indicated by point C. The dot-dash and dashed curves are examples of relatively high quality and lesser quality beams, respectively. The labels A, B, etc. are used for reducing these curves to single-value BQ parameters as discussed in the text. Point P is an example of a programmatic goal, in which case the goal is met if the measured PIB curve passes through or above this point. In this example, note that the "ideal" curve does not have the largest PIB at large bucket radii. This can occur whenever the ideal beam is a flat top (leading to a substantial level high-angle diffraction) and when the beam it's compared to is a truncated Gaussian, with both beams being forced through the same near field aperture. The dot-dash curve is indicative of a truncated Gaussian intensity distribution with low spatial frequency phase errors, but not shown is a low level of diffraction modulation that would actually be present.

Programmatic BQ requirement: Requirements can be stated in any number of ways. One example is that a point on the PIB map that the laser must exceed is specified. For example, the beam must contain at least 70% of the total power within a bucket radius of $2\lambda/D$. Such a requirement is neither a vertical nor a horizontal specification.

3.2.1. Power in the bucket curve. An example of the PIB curve is shown in Fig. 3, which shows three example curves, each of which is simply the result of integrating the far-field irradiance distribution over a solid angle whose radius varies according to the abscissa value. A natural unit for that axis is the size of the central diffraction lobe, but by convention we often instead use λ/D, where λ is the wavelength and D is some characteristic dimension of the near-field beam. If the beam is round and D is the near-field diameter, then the radius of the far-field central lobe is $1.22\ \lambda/D$ when the near field is a uniformly illuminated disk with uniform phase. An example ideal curve (i.e., diffraction limited) is given by the solid line. This reference curve is generally calculated from an *ideal near field.*

3.2.1.1. Specification of the ideal near field. Because the goodness of a beam depends on comparison of a measured PIB to an ideal PIB, with the latter being calculated from an ideal near field, the choice of this ideal near field is important. Actually, it is the shape of the ideal near field, not it size, that matters because we will compare the measured far-field pattern to the ideal far-field pattern—with both near fields scaled to the same size. Similarly, it makes no difference conceptually whether we consider the near field to be measured at the laser or instead at the output of the beam director.

Specification of the near-field shape will affect the measured BQ directly, and this shape might be chosen differently for different purposes. Here are some examples:

- Consider a segmented output beam (i.e., side-by-side apertures with dead space between). The laser physicist may be interested primarily in phase errors of the near field, without regard to how the segmentation affects the far-field pattern. So he may define the ideal near field to include the segmentation explicitly, in which case he will measure BQ ≈ 1 when the phase errors are eliminated.
- The systems engineer may desire a beam with central obscuration for use with a Cassegrain-telescope beam director, and hence he may define the ideal near field as having an obscuration. Nominally, the system engineer will be maximizing target lethality.
- The program manager may need a simply defined BQ value for tracking progress or ball-parking the laser capability, and he may choose the ideal field to be a uniformly illuminated disk that encircles the laser output aperture, whatever its actual shape.

3.2.1.2. Specification of the far-field bucket shape. The far-field bucket is not necessarily round. Following from the above, the ideal near field might be chosen to have a nonunity aspect ratio (e.g., to match the shape of a slab gain media). So in such cases the ideal far-field central lobe will not be round, and the shape of the far-field bucket would be adjusted accordingly. Otherwise, the measured PIB curve responds differently to BQ errors on one axis as compared to equivalent errors on the other axis.

3.2.2. Conversion of PIB data to single-value BQ numbers. With the ideal near field specified and the far-field bucket specified, one can then calculate the reference PIB and measure the actual PIB. Whereas the measured PIB constitutes the key characterization data, we now discuss conversion of the PIB to a single number. This may be useful for programmatic purposes, even though information is lost.

3.2.2.1. Horizontal beam quality measurement method. Referring in Fig. 3 to a measured beam corresponding to the dashed line, the *horizontal beam quality* is defined as the ratio of segment DF to segment DE, specifically,

$$\mathrm{BQ}_{\mathrm{horizontal}} = \frac{\text{actual radius at given power}}{\text{ideal radius at same power}}. \tag{1}$$

Recalling the first two paragraphs in Sec. 3.1, this measure scales as the ratio of two quantities that are each an angle $\times$ length product. In Fig. 3 the fraction of power is set arbitrarily at 75%, although there is no standard choice for this value. If a small value is chosen, e.g., 30%, then our BQ metric gives no information regarding where 70% of the power lies. If a large value is chosen, e.g., 95%, then we know what solid angle contains nearly all the power but nothing of how it is distributed.

The horizontal definition is useful for beams of any quality, particularly if the distribution function is roughly known, e.g., Gaussian.

3.2.2.2. Vertical beam quality measurement method. Referring now in Fig. 3 to a measured beam corresponding to the dot-dash line, the *vertical beam quality* is defined using the ratio of the line segments AC and AB, specifically,

$$\mathrm{BQ}_{\mathrm{vertical}} = \left(\frac{\text{ideal power in bucket}}{\text{measured power in same bucket}} \right)^{1/2}. \tag{2}$$

Recalling the first two paragraphs in Sec. 3.1, this measure scales as the ratio of two quantities that are each an angle $\times$ length product. Often, the bucket is chosen to be the size of the ideal far-field central diffraction lobe. More exactly and in that case, it is defined by the calculated line of zero intensity surrounding the lobe. As a detail, when the near field has limited spatial symmetry, there may be no precisely defined line of zero intensity. For a given beam and these definitions, it is generally true that $\mathrm{BQ}_{\mathrm{vertical}} \neq \mathrm{BQ}_{\mathrm{horizontal}}$ and one cannot be converted to the other without knowledge of the far-field distribution function.

If the vertical beam quality is measured for the beam corresponding to the *green* line in the figure, then the fraction of PIB (i.e., at 1.22 λ/D) is only about 15%. Therefore for this beam, if we have knowledge of only the vertical BQ (but not the entire PIB curve), we would know that 15% of the power lies within 1.22 λ/D, but the whereabouts of the other 85% is unknown.

Accordingly, the vertical beam quality definition is of limited value for determination of lethality for beams that are not close to DL. If we restrict use of this metric to cases in which $\geq$50% of the power falls within the bucket, then the region of applicability is that $\mathrm{BQ} \leq 1.3$.

A two-parameter description of the beam will carry more information than a single parameter (e.g., one could report both vertical and horizontal BQ), and of course the full PIB curve is more useful still. In some cases the far-field irradiance distribution can be reasonably well described by two Gaussian beams, a DL core and a wider pedestal (private communication, Butch Deuto, Air Force Research Laboratory). The partition of energy between these two components would depend on the particular laser.

3.2.3. The *two-hard-aperture* measurement approach: An implementation of vertical beam quality. In the case that the actual near-field beam has hard edges, e.g., a top-hat intensity distribution, a particularly simple and robust way to make a BQ measurement to it is as follows:

(a) Specify the ideal near field to be of some given shape with hard edges and with uniform illumination and phase. Keep in mind that from a systems perspective, the near field will nominally be specified so that it fully fills the beam director and that matching the laser to the beam director could involve beam reshaping.
(b) "Measure" the near-field size of the actual beam by forcing it through this aperture. This can be done with a low-power sample of the beam and/or in some cases with a camera-based virtual aperture as discussed later.
(c) Calculate a far-field bucket size based on the given near field, and measure the fraction of total power that is captured by this bucket as an indication of goodness. If the fraction

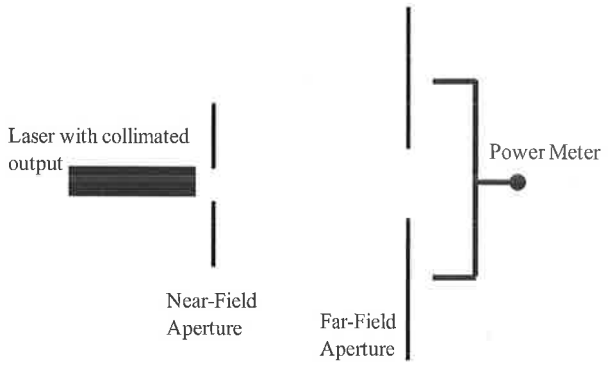

Fig. 4. Geometry of the two-hard-aperture BQ measurement. This measurement is nominally based on the ratio of measured power, with both apertures present, to the measured power with both apertures absent. See the text for discussion of power intercepted by the near-field aperture and replacement of both apertures by cameras. The figure is not to scale; the far-field aperture must satisfy the condition $(D_{\text{near-field}})^2 \ll L\lambda$, where D refers to the near-field aperture diameter, L is the distance between apertures, and λ is the wavelength. A common choice would be that $D_{\text{far-field}} = 2.44\ L\lambda/D_{\text{near-field}}$, which is the far-field diameter of the central lobe from a near-field circular flat top.

of total power intercepted by the near-field aperture is negligible, then this is a measure of beam quality. (As a detail, the general use of the term "beam quality" refers to properties of photons emitted from a near-field aperture. But in this section we are considering the more general case in which some photons may be intercepted by the near-field aperture.) Referring to the first two paragraphs in Sec. 3.1, this measure scales as the ratio of two quantities that are each a solid angle × area product. If this fraction is inverted and the square root is taken, then it scales equivalently to Eq. (2), which is the more conventional BQ terminology. Of course the BQ derived in this manner refers only to the photons that were transmitted through the near-field aperture in order that the measurement is self-consistent.

Given that both the near- and far-field measurements can be made with a hard aperture, we will call this the *two-hard-aperture* approach, and its geometry is depicted in Fig. 4.

With this, we have completely sidestepped the complication of measuring the actual near-field intensity distribution and deciding where the edge is. This is a considerable simplification, and one justification to this approach is that in a DE system, we are essentially making a two-hard-aperture measurement when we engage a target through a fixed-size beam director and seek a limited spot size at the target that is based on lethality considerations.

Now consider an important point regarding interception of the actual beam by the near-field aperture. This aperture must closely fit the beam and not be overly large; otherwise the measured BQ will inadvertently be worse than the actual BQ. This is because a near-field aperture that generously exceeds the beam size will result in a calculated ideal far-field bucket that is unnecessarily small. Accordingly, to make a practical BQ measurement with a given near-field beam, the size of the near-field aperture would be reduced until it just begins to intercept the beam.

Considering that we need a self-consistent measure of power and beam quality, that our BQ measure refers to power downstream of the near-field aperture (Fig. 4), and that the near-field aperture can intercept power, then there is in general a trade-off between the measured power and the measured BQ. This is because as the near-field aperture size is reduced, the measured power falls while the (calculated) far-field bucket is enlarged.

The approach to this trade that worked well in the JHPSSL program was that the independent measurement team allowed each contractor to choose the near-field aperture size. That is, the contractor made the trade between power and BQ as they saw fit according to the program goals. Recall that the *shape* of the near field would have been predetermined at the program outset. The JHPSSL beams had relatively hard edges, so that this trade-off was straightforward.

Implementation issues and consideration of beams with very soft edges, e.g., Gaussian beams, is discussed in Sec. 3.3.1.3.

3.3. Power in the bucket measurement and hardware details

Figure 5 depicts the generic hardware components useful for HEL characterization. Beginning in the upper-left-hand corner, the HEL beam is attenuated by a known amount and then sent to the diagnostics station. The near-field beam may be intercepted slightly by the near-field aperture and can alternatively be examined by a virtual aperture (i.e., a nonintercepting aperture). A telescope reimages this aperture (real or virtual) to a remote location if needed. The absolute power meter is shown. A transform lens forms a far-field image at the locations of the far-field camera, far-field pinhole, and lateral position sensor. The latter provides a high-bandwidth determination of the jitter power spectral distribution (PSD), from which the rms jitter value can be calculated. The fast steering mirror (FSM) shown may be needed if pinhole-based BQ measurements are to be separated from the effect of jitter. This mirror can be commanded by signals from the lateral position sensor so that the pinhole tracks the beam position. A similar system was used for the JHPSSL

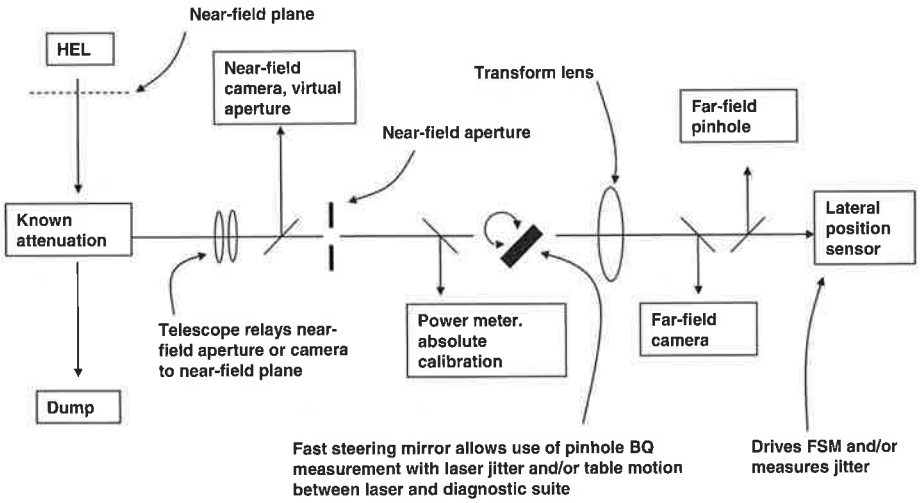

Fig. 5. Generic hardware configuration for HEL characterization that permits simultaneous measurement of laser parameters.

independent measurements.[1] The generic system of the figure represents an essentially complete diagnostic except for polarization analysis.

Details and hardware issues are discussed in the following subsections.

3.3.1. The near-field aperture. There are a number of implementation details to understand with the near-field aperture. First, note in Fig. 5 that the aperture is placed at a low power point so that heat loading is not an issue in the real aperture case. This aperture nominally lies in the plane that has been designated as the output plane of the laser, and this plane will not necessarily lie within the physical bounds of the diagnostics hardware. Accordingly, the telescope in Fig. 5 projects the near-field aperture to an agreed-upon position.

3.3.1.1. Aperture sizing and alignment in the case of a real aperture. The implementation is conceptually straightforward in the case of top-hat beams. But there are clear challenges in practice, namely alignment and sizing. Recall from the above that fitting the beam through this aperture is essentially a trade-off between beam quality and power and that this trade will be made to best suit the program specifications (e.g., required power and beam quality). For alignment, it is convenient to be able to view the beam strike from the input side of the aperture. For sizing, ideally the aperture is adjustable. This is workable for simple beam shape but nearly unworkable for complex situations as would be the case with segmented beams. The alternative of a virtual aperture is considered next.

3.3.1.2. Use of a virtual near-field aperture. If a complicated shape is chosen for the ideal near field, e.g., one carrying all the details of a segmented beam including the dead space, then the use of a physical near-field aperture becomes impractical due to the problem of getting correct alignment, correct size, and the desired trade-off between BQ and power—all at the same time.

In this case a *virtual near-field aperture* could be considered (see Fig. 5). In general this would be a camera-based aperture, projected to the desired near-field plane. This virtual aperture can be any shape, and the power falling outside the aperture can be computed from the image. In Sec. 3.3.4 there is discussion of the problems of getting accurate power integrals from cameras, and that discussion can also be applied to the near-field camera as one attempts to determine how much power falls outside the boundary of the aperture.

Assuming that the camera works perfectly so that the power integrals are accurate, the problem with this approach is that the far-field pattern is affected by the photons that were outside the virtual aperture; hence the measurements of power (i.e., photons falling inside the virtual aperture) and BQ are not self-consistent. Therefore, before using this approach one should establish with modeling that the effect on the BQ measurement is acceptable. As an example, in the JHPSSL program, when a virtual near-field aperture was used, it was required that this aperture be large enough to capture at least 95% of the total power. Also, the power falling outside this aperture was subtracted from the reported power, i.e., the power taken to be (essentially) consistent with the reported BQ.

We stress that a physical aperture is the gold standard for determination of near-field size and any other means of measurement is ideally validated/calibrated against a physical aperture.

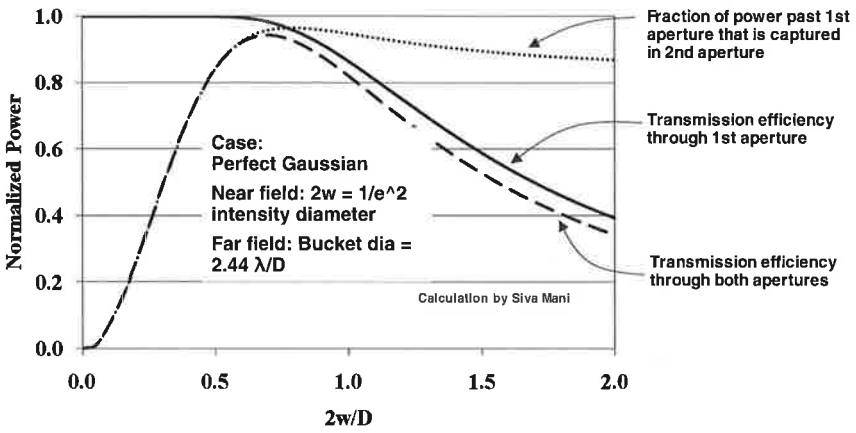

Fig. 6. Power lost as near-field beam size is varied while BQ measurement apertures are held fixed. Refer also to Fig. 4. See the text for explanation.

3.3.1.3. Use of near-field soft-edge beams in a two-hard-aperture measurement. Up to this point we have considered the case in which the actual beam has a near-top-hat distribution, so that the near-field aperture intercepts only a small fraction of the laser power while providing a good measure of the beam size.

Now consider the case in which the actual near-field beam has a soft edge, e.g., a Gaussian intensity distribution (but we still force it through a hard-edge aperture), so that the trade between apparent power and apparent beam quality becomes an important consideration (discussed in Sec. 3.2.3). The issue is that, with a soft-edge, near-field beam, there is a continuum of choices between measured power and measured BQ, so that the meaning of these measurements is not clear.

In this case and from a systems viewpoint, i.e., delivery of power through a beam director to a target, measuring separate power and BQ values may have limited usefulness. A better metric is the fraction of total laser power (i.e., prior to any aperture interception) that falls in a specific far-field spot, regardless of how much power is lost to the near-field aperture. This is what the two-hard-aperture approach measures. But in this case of the soft-edge near field, it is not a measure of *beam quality* but instead a measure of *efficiency to bucket.* An equivalent term, power in the bucket efficiency (PIBE), is also in use.

If that is a candidate metric, we can ask whether maximization of this metric leads to nonsensically large interception by the near-field aperture, i.e., nonsensical in the sense that it would have little relationship to any realistic system.

To explore this, recall that Fig. 4 shows the geometry of interest and now examine the calculation (unpublished calculation, Siva Mani, Schafer Corp.) shown in Fig. 6.

In Fig. 6 we consider that a two-hard-aperture measurement is being made and that the apertures are fixed in space (in both the distance of separation and the two diameters) such that the far-field aperture exactly matches the Airy disk of the near-field aperture. Furthermore, the size of the input put beam is varied, and in Fig. 6 we calculate the beam interception by the two apertures. The near-field intensity distribution is Gaussian with $2w$ defining the diameter of the e^{-2} point in intensity. The input beam is diffraction limited in the sense that it has a flat phase profile, and the far-field bucket diameter is fixed at $2.44\lambda/D$. The plotted lines show the power intercepted by the two apertures and also the total transmission efficiency.

Our *efficiency to bucket* metric relates to the dashed line, and we wish to see what fraction of power is intercepted by the near-field aperture when this metric is optimized. The maximum metric value (dashed line) of 0.94 occurs at $2w/D = 0.7$, and for this condition 2% of the power is lost to the near-field aperture (solid line). If that seems high for a HEL system, how far from the metric optimum would we be for lower interception? At $2w/D = 0.6$ there is 0.4% interception and the metric falls to 0.92. At $2w/D = 0.55$ there is 0.1% interception and the metric falls to 0.89. So, for these assumptions, nearly all the near-field aperture interception can be eliminated while still maintaining near-optimum efficiency to bucket. Therefore we conclude that the *efficiency to bucket* metric, using two hard apertures, can be useful even for Gaussian beams. Indeed, this will be one of the primary metrics for the RELI program.

3.3.2. Relative placement of the output-plane, near-field aperture, transform lens, and pinhole. Given that the designated laser output plane, nominally a point of collimation and also sometimes a point with well-defined beam edges, may lie inside the laser and far from the diagnostic hardware, we consider now basic rules for placement of key hardware items.

Bear in mind that the nominal setup for the BQ measurement is that the laser output is perfectly collimated at an agreed-upon near-field plane and that this plane is colocated with the near-field aperture and transform lens. The far-field pinhole (or camera) is located at a distance f behind the lens, where f is the lens focal length.

The geometry of the above paragraph may not be realized due to limitations of the diagnostic hardware or due to physical separation of the laser and diagnostics. There are two basic issues. One is that the beam may change size between the near-field aperture and transform lens due to imperfect collimation, and the other is diffractive effects as the beam propagates.

Regarding diffractive effects, these can be safely ignored only if the near-field plane, the nearfield aperture, and the transform lens are all "in the near field of each other." The criteria for that can be stated in terms of an effective Fresnel number, $F_{\text{eff}} \equiv {a_{\text{eff}}}^2/(L\lambda) \gg 1$, where L is the separation distance of interest and λ is the wavelength. Parameter a_{eff} is the length scale of the highest spatial frequency error of interest. It could be much smaller than the beam diameter, particularly considering that a large-aperture HEL can exhibit unwanted, small-scale structure from a variety of sources including cooling nonuniformities, actuator print-through on deformable mirrors, or intensity variation stemming from the dead space associated with beam combining via aperture tiling.

We now consider several off-nominal conditions in more detail.

3.3.2.1. When the near-field plane is not close to (and is in front of) the near-field aperture. This refers to the case in which the effective Fresnel number between the designated near-field plane and near-field aperture is too low. It could happen in systems that do not have projection capability for the near-field aperture. In this case a physical optics model can be used to estimate whether the BQ measurement is compromised. For large beams ($\gg$1 cm) and modest distances ($\ll$10 m) the projection telescope is generally not necessary.

3.3.2.2. When the near-field aperture is not close to (and is in front of) the transform lens. This could happen in the case in which the near-field aperture is projected outside the

diagnostic hardware to the near-field plane. There are two cases to consider here:

(a) If the beam size at the near-field aperture does not match the size at the transform lens due to lack of collimation, this is acceptable only if it is compensated for in the BQ calculation.
(b) If the effective Fresnel number between the near-field aperture and the transform lens is too low; see Sec. 3.3.2.3.

3.3.2.3. Placement of far-field pinhole relative to transform lens. Consider the general question of how to place the far-field pinhole (or camera) relative to the transform lens. Consider further that the near field is colocated with the near-field aperture. When the near field is close to the lens and perfectly collimated, then the Fourier plane (of the near field) is found exactly a distance f behind the lens, where f is the focal length. Now, keeping the near-field plane and aperture together, consider increasing their distance from the transform lens by an arbitrary amount such that the condition $a^2/(L\lambda) \gg 1$ is no longer satisfied, where a is the beam radius. (For simplification, consider also that the transverse scale length of aberrations of interest is also of order a.) It turns out that in this case the Fourier plane is still located exactly distance f behind the lens,[2] but now there is a position of best focus (i.e., a slightly smaller spot) at a distance slightly farther from the lens. This is due to diffractive curvature of the wavefront upstream of the lens.

A question to ask is whether these two planes, the Fourier plane and the plane of best focus, are indistinguishable. More specifically, we ask whether the focal depth of field encompasses both planes. The depth of focus is given roughly by $\lambda f^2/D^2$, where f is the focal length and D is the beam diameter at the lens. The offset between the two planes is maximized when curvature due to diffraction is maximized, and this occurs when the Fresnel number between the near field and lens approaches unity, roughly, and that maximum offset is about equal to depth of focus. So, the two planes are essentially indistinguishable so long as the Fresnel number between the near field and lens exceeds about 10. This is a pleasing result, because the position of best focus is easy to find and the Fresnel number requirement is generally easy to satisfy. If the BQ measurement procedure is to work at the position of best focus, then small focal errors at the near-field plane are unimportant. In fact, with small focal errors there would be no simple way to find the precise location of the Fourier plane anyway. Our conclusion is that we should use the position of best focus.

3.3.3. Avoiding need for absolutely calibrated detectors. The PIB BQ measurement method requires that the ratio of the PIB to the total power be measured. That is, the PIB BQ is really based on a ratio, not a power. One way to find this ratio is to measure the true power falling in the bucket on an absolute basis and to also measure the total laser power on an absolute basis. If this involves two different detectors and two different optical paths that may have different fields of view or polarization sensitivity, then this strategy may lead to problems of accuracy and reliability.

An alternate method is to use a linear detector, not requiring absolute calibration, that can trace out the entire PIB curve as in Fig. 3. If this curve extends to sufficiently large bucket size that nearly all photons are captured, then it is already calibrated to unity power. So the requirement to avoid the need for an absolutely calibrated detector is that the PIB detector has a sufficient field of view to capture the entire beam. If the laser power varies in time, it may be advantageous to use a separate detector for normalization against total power.

In principle, one can dispense with the full PIB *curve* and use just two far-field pinholes, provided that we know that the larger pinhole captures all the power and that the smaller is the size of interest, nominally around 1 λ/D. A drawback to this is that there is no a priori way to know how large the angular acceptance of the larger pinhole needs to be. A necessary, but not sufficient, indication that the acceptance is large enough comes from measuring the whole PIB curve and observing that the slope does indeed tend to zero in the upper-right-hand corner. But this indicator can be misleading for at least two reasons:

(a) In the case that a camera is used to measure the PIB curve, there are potential problems, which are discussed in the next section
(b) There are cases in which the laser may produce a DL core but at the same time also produces substantial additional energy at high angles. A laser employing beam combination via a diffractive optical element (DOE) would display this behavior if the beam phases are not adjusted properly.

We stress the importance of demonstrating that the PIB values are actually normalized as close to 100% power as possible. *It is important that the far-field angular acceptance of the diagnostic system exceeds the far-field angular content of the laser.* (Of course any hard-edge beam has, by necessity, a small fraction of power emitted at very high angles, so we are speaking here about what can be done realistically.) Each sensor in a complex optical system will in general have a different field of view, in both position and angle space. This must be accounted for, because the laser may have an unknown angular content.

3.3.4. Potential Difficulty with Camera-Based PIB Curves – The Value of Pinhole Measurements. Generally, the far-field PIB curve, calibrated as described in Section 3.3.3., can be generated either with a single detector (e.g. photodiode) and a series of pinholes, or from a camera image using integration in software. It was the experience of the independent measurement team for the JHPSSL program, that silicon-based CCD detectors (at lease those tried) may not have sufficient accuracy to generate reliable PIB curves. An indication of the problem was that apparently spurious photons were detected at large angles, so that the PIB curves did not asymptote as expected. Others have observed similar behavior (private communication, Ken Widen of Boeing, and Stu McNaught of Northrop Grumman).

The specific problem with the camera is that to capture all the energy, one may need to integrate the image to high angles (e.g. $\gg 10\lambda/D$). But over much of this integration area the signal level may be low, in which case camera noise can dominate and the PIB curve would be unreliable. One solution is to eliminate the camera entirely and rely on pinholes. (The detector behind a pinhole, e.g. a photodiode, can have a large dynamic range, giving in essence an accurate area integral.) But there are potential drawbacks to an all-pinhole approach:

(a) It can be difficult to ensure that the pinhole is well aligned to the beam. If jitter and beam pointing drift are present, alignment may be impossible without servo-driven tacking. (A far-field beam diameter of 100 microns would not be unusual.)
(b) Many pinholes would be required to trace out the whole PIB curve.
(c) The far-field bucket shape of interest may not be round, in which case the camera is quite versatile by comparison.

So one approach is to combine the camera and pinholes using a basic strategy of limiting the area over which the camera must integrate. An example of one approach to this is shown in Fig. 7, which shows a PIB curve for a situation in which most of the energy is accounted

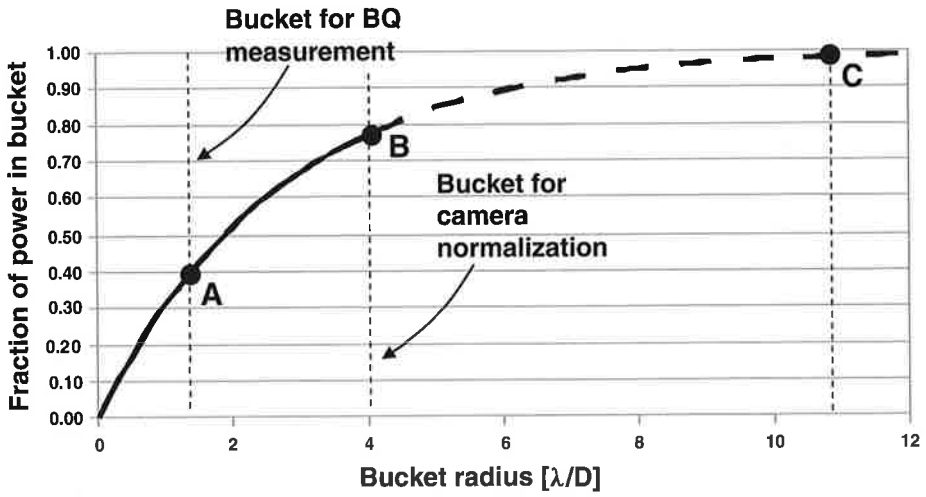

Fig. 7. Use of camera and pinholes together to provide a correctly normalized PIB curve. The ratio A/B is measured by camera, and the ratio B/C is measured by pinholes. Assuming that C represents essentially all of the beam (which must be verified), then point A is located properly as (A/B)×(B/C), which in this case is ≈0.40. Accuracy at point A diminishes if the radius at B is much greater than the radius at A, and of course no information is available for radii between B and C. An approach using pinholes both smaller and larger than point A is discussed in the text. See discussion in the text regarding jitter and alignment.

for within 4 λ/D, and the bucket for the goodness measurement is arbitrarily chosen for this example to be 1.5 λ/D. The camera curve at 4 λ/D is normalized with a pinhole of exactly that size. There are actually two pinholes, one is 4 λ/D, and the other is large enough to capture all photons. With this arrangement, the camera integral covers comparatively little area of low signal and the accuracy of the PIB curve improves. The camera can also be used to compute PIB for non-circular buckets. *It is important that the alignment of the pinhole with respect to the camera frame be established with precision.*

If jitter is not a problem and alignment between the pinholes and beam can be assured for small buckets, an approach improving on that of Fig. 7 is to use a series of pinholes that span the 1.5 λ/D range (for this example) and use the camera to interpolate between these points, or for non-circular buckets. It should be emphasized that this places more emphasis on correct alignment than does the 4 λ/D example above.

Given the need for measurement methods that will give equivalent BQ values when different teams measure the same beam with different hardware, we stress the point that the pinhole method is the gold standard for PIB curves. Camera-based methods should be validated/calibrated against pinholes.

3.3.5. Pointing jitter. For pointing jitter, we are considering the situation in which the beam is spatially stable at the plane defining the near field (Fig. 5) but that it may not be stable in angle. Therefore the effect of jitter is movement of the far-field spot, and if time averaged, any jitter is equivalent to an apparent worsening of the beam quality. Also, separation of the laser and diagnostics (they may be on separate tables) can lead to additional apparent jitter.

Whether jitter is to be considered a component of the BQ, or instead to be separately reported, depends on program specifications. If they are separated, then the term *instantaneous BQ* may be useful for aberrations other than pointing errors. It all depends on program needs; in the JHPSSL program both the dynamic pointing error and static focal errors were considered not to be part of BQ, and in the RELI program the plan is to include high-temporal-frequency pointing errors in the BQ, but not low-frequency errors.

Certainly a statically aligned pinhole can be used for measurement of average BQ (i.e., including jitter) but is problematic if jitter is to be reported separately from (instantaneous) BQ. However, and considering the relatively large 4 λ/D normalization example of the preceding subsection, if the jitter is small compared to the size of the normalization pinhole, then a fast camera can be both normalized with this pinhole and used to separate jitter from BQ. An exposure time of Δt freezes all jitter components at frequency $f \approx 1/(2\pi\,\Delta t)$ and lower. The BQ can be measured on a frame-by-frame basis, centering the bucket on the beam in each frame. These BQ values can be averaged as needed, and the rms jitter value (but not the jitter PSD) is also available from the same data. This procedure was used in the JHPSSSL program. *Note that with the CCDs examined in that program, the camera integration (to calculate the PIB curve) became less reliable as the exposure time was shortened.*

Another technique used with the JHPSSL hardware was to use adaptive optics (AO) to make the pinhole track the laser beam. Figure 5 shows a FSM and lateral position sensor for such a tracking loop. It may also be possible to use an AO system already present in the HEL to provide this function for the BQ measurement.

Whether a fast camera is freezing out jitter or the pinhole tracks the beam centroid, there may be high-frequency jitter components that cannot be separated from BQ. Lumping the high-frequency jitter components together with the usual BQ aberrations (i.e., terms higher than tilt and focus) is reasonable from the perspective of the overall system (HEL plus beam director), because there is likely to be a FSM present and it will have limited bandwidth. This viewpoint is taken in the RELI program.

The ideal reporting of jitter is a power spectral density (PSD) distribution function. This can be reliably measured from dc to high-frequency components by lateral position sensors (Fig. 5).

4. Other Characterization Parameters

4.1. Determination of turn-on time and pulsing capability

Bringing the laser from a full-off to a full-on condition requires transitioning through various levels of readiness, each of which is associated with some minimum time interval. The time span from full-off to full-on is the sum of these transition intervals. The rate of consumption of platform resources for each readiness level is important for mission utility analysis.

This subsection is written under the assumption that there exists a key transition from a *ready state* to a *fire state* that can be measured by observing the laser output when starting from a nonlasing condition. It is also assumed that this ready state is defined by some suitably low electrical power draw from a mission standpoint and that the system is capable of dwelling in this ready condition. The whole notion of ready states and turn-on time must be considered against the mission, platform, and critical resources. (For example, if the critical resource is a cryogen supply, considerations different from those here may be appropriate.)

The laser can be characterized by specifying the minimum time associated with each level of readiness. A graphical representation of the turn-on process is shown in Fig. 8

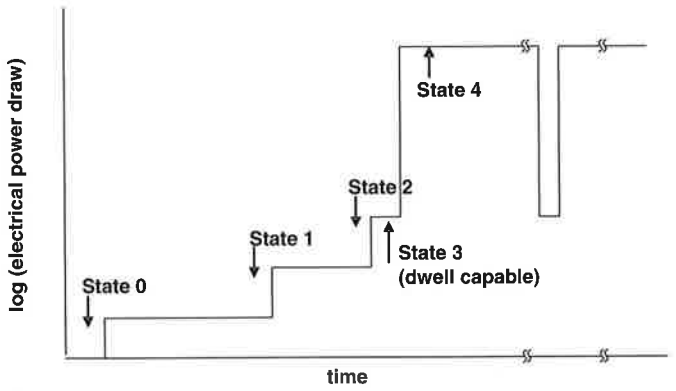

Fig. 8. States of operations relevant to the determination of start-up time. Of primary interest is the ready state (state 3) and for HEL fire state (state 4).

based on the electrical draw as a key discriminant. State 0 is the cold start, and state 4 is full laser power. States 1, 2, and 3 are arbitrary states of varying readiness. They could pertain, for example, to different degrees of thermal conditioning among the optical components and cooling subsystems.

Of particular interest will be state 3 and its associated power draw. This is the ready state, from which the laser can be fired *immediately*, i.e., of order 1 s or less. The turn-on time is then measured as the transition from state 3 to state 4. Each program will need to determine the appropriate definition of turn-on time and allowed time increment, taking into consideration that equilibration times in each of the states (Fig. 8) are important. An example definition is that turn-on time is the time, starting from state 3, to achieve the program-specified power simultaneously with the specified beam quality, polarization state, and jitter. It may happen that a laser produces full power prior to the time that the BQ has optimized.

Nominally some pulsing capability would also be needed, and this requirement must also be determined by the program. The requirement might be to demonstrate a particular pulse train or demonstrate a maximum allowed cycle time. If the system is repeatedly driven from ready state to fire state as fast as possible, then a cycle time can be measured (i.e., the spacing between the leading edge of two successive pulses).

4.2. Efficiency

Assuming the case of an electrically driven laser, below are some useful terms relating to efficiency, along with suggested definitions:

Electrical efficiency: A general term that can be applied to any group of components.

Optical-to-optical efficiency: Conversion efficiency from pump power to laser power. This term always needs detailed specification, particularly relating to how the pump power is determined (e.g., total diode emission, diode power after diode combining, power incident on gain media, or power absorbed by the gain media).

Electro-optical efficiency: A narrowly defined term for optical power produced divided by electrical power supplied to drive the gain media. This term is most useful when listed in conjunction with a given output power, beam quality, and explanation of at what point in the system the electrical power draw is measured (e.g., before or after power conditioning). If

Table 1. Generic characterization report template (it is assumed that a requirement has been identified for each parameter)

ID	Parameter	Deliverable data	Comments
1	Power, run time	Graph: Power versus time. May require evaluation of polarization state.	Should be confirmed with independent measurement. See Sec. 2.
2	Turn-on time	Graph: Power versus time. An agreed-to ready state should be identified, and power draw in ready state should be indicated.	Detailed definition required. See Sec. 4.1.
3	Pulse capability	Graph: Power versus time for specified pulse train; alternate: provide graph showing cycle time.	Detailed definition required. See Sec. 4.1.
4	PIB curve	Graph: Measured and ideal curves.	Near- and far-field sizes should be measured with either physical apertures or detectors calibrated to physical apertures. See Sec. 3.
5	PIB-based BQ	Report single value along with simultaneous beam power. Need identification of near-field plane, near-field shape, and size/shape of far-field bucket.	Same as above. Include both BQ-vert and BQ-horiz if BQ-vert > 1.3.
6	Jitter	Report PSD curve and rms values in each plane.	See Sec. 3.3.5.
7	Electro-optical efficiency	Report single value at given laser power and BQ.	Detailed definition required. See Sec. 4.2.
8	Electrical efficiency to bucket	Report single value at given laser power and BQ.	Detailed definition required. See Sec. 4.2.
9	Wall plug efficiency	Engineering analysis report.	Need to specify interface to platform.
10	System size and weight	Engineering analysis report.	Platform specific.

an oscillator-amplifier configuration is used, one would specify whether low-power stages are included.

Electrical efficiency to bucket: Similar to electro-optical efficiency, except that only the power falling in a specified far-field bucket is counted. One would need to specify whether the electrical draw is only that for optical pumping or whether more system elements are included.

Wall plug efficiency: The *overall* system efficiency including *all* ancillary systems. Requires careful definition if some of the ancillary functions are provided by the platform. In a laboratory setting, this efficiency would generally be determined by engineering analysis so that, for example, inefficient house cooling systems do not degrade the result.

Except for the optical pump power and laser output power, measurement of these parameters should be straightforward using standard electrical test instrumentation that is subject to a periodic calibration protocol. An explanation of how/where each measurement is made must be included in any test report.

4.3. Other optical properties

Some optical properties that are potentially relevant to DE applications are not discussed in this document. These include the following:

(a) Polarization state: May affect design of optical systems used in conjunction with the HEL, e.g., a beam director and pointing/tracking system.
(b) Spectral content (wavelength and width): May affect atmospheric propagation and target coupling.
(c) Short-pulse effects: May affect atmospheric propagation and target effects due to nonlinear interactions.

5. Example Outline of Laser Characterization Report

Table 1 is a generic report template identifying useful characterization parameters, written in the context of a contract deliverables list. It includes references to the preceding sections, as well as some straightforward items not already discussed. Size and weight are called out, but these are not developed in this paper.

6. Acknowledgment

This work has been supported by the HEL-JTO under contract FA945105F0239.

References

[1]Edwards, B., S. Di Cecca, D. Murphy, and J. Slater, "JHPSSL Government Diagnostics System," presented at Tenth Annual Directed Energy Symposium, Huntsville, AL (2007).

[2]Goodman, J. W., *Introduction to Fourier Optics*, McGraw–Hill, San Francisco, CA (1968).

[3]"JHPSSL Technical Notes for Laser Characterization," JHPSSL Technical Evaluation Committee (2004).

[4]Lawrence, R. C., "Active Wavefront Correction in Laser Interferometric Gravitational Wave Detectors," Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA (2003).

[5]Ross, T. S., and W. P. Latham, J. Directed Energy, **2**, 22 (2006).

[6]Siegman, A. E., "How to (Maybe) Measure Laser Beam Quality," Tutorial presentation at the Optical Society of America Annual Meeting, 11 April 2011, Long Beach, CA (1977). Available at http://www.stanford.edu/~siegman/beams_and_resonators/beam_quality_tutorial_osa.pdf.

The Author

Dr. Jack Slater joined Schafer Corporation in 2002 and presently leads the Schafer team supporting the High Energy Laser Joint Technology Office (HEL-JTO). He has been closely connected with the JTO's 100-kW solid-state laser program, including the Government-sponsored independent testing of high-power lasers. He received his Ph.D. in atomic physics from the University of Colorado and worked previously at STI Optronics, with emphasis on free-electron lasers, and with the physics group at the Idaho National Laboratory. He is a Fellow of the Directed Energy Professional Society.