Journal of Physics: Conference Series **299** (2011) 012014 doi:10.1088/1742-6596/299/1/012014

The Jefferson Lab Free Electron Laser Program

Stephen V. Benson, David Douglas, George R. Neil, and Michelle D. Shinn

Jefferson Lab, Newport News, Virginia 23606 USA

Abstract. The development of superconducting radio frequency (SRF) accelerator technology for CEBAF's nuclear physics program set the stage for its application to a number of other efforts. We describe below the development of the a major advance in Free Electron Laser (FEL) Technology based on SRF linacs. The Jefferson Lab efforts achieved three orders of magnitude increase in the power delivered by FELs and firmly established the viability of energy recovering linac technology. We describe the details of the physics and engineering challenges addressed to accomplish this effort and then discuss some of the applications performed using the light source. We conclude with a look at the planned directions for the program in the future.

1. Introduction

The development of reliable and cost effective SRF accelerators as demonstrated by the operation of the CEBAF accelerator in 1994 [1] set the stage for an application to high average power Free Electron Lasers. It had been realized early in CEBAF construction that the system could be utilized for an FEL [2] but it wasn't clear in the early stage of this study what the specific advantage would be over other systems or what the goals of such a program might be. FELs had been considered as candidates for high average power almost from their inception and the decade that followed was a wild ride of great hopes driven by large infusions of cash from Strategic Defense Initiative (SDI) interest. Ultimately these efforts in the US lapsed into a number of low level applied research studies with little to show from the nearly $1B investment [3]. At that time there was continuing improvement in the performance of solid state lasers with optical parametric amplifiers extending to longer wavelengths and beginning to answer needs for laboratory sources of tunable coherent infrared radiation. Yet there was a need for capabilities that could not be satisfied by existing or near-term foreseeable lasers: short pulses of high energy in high-repetition-rate pulse trains, tunable throughout the infrared region. These needs spanned basic research in condensed matter, atomic and molecular physics and, if powers up to 100 kW could be produced, industrial applications. Support for such an effort was eventually obtained through the Office of Naval Research establishing a development effort for a defense system against cruise missiles that continues today. Initial studies [4] considered use of the CEBAF accelerator but the fear of compromise of the nuclear physics program led to the construction of a dedicated FEL facility utilizing not only SRF technology but applying energy recovery [5] for improvement of performance, and efficient, cost-effective scaling to higher powers.

The article below summarizes the successful development of a superconducting, energy-recovering Free Electron Laser which became the highest power FEL in the world. It also established a new technology, Energy Recovering Linacs (ERLs), as a basis for development of research light sources for the 21st century and other applications. We follow the physics

Journal of Physics: Conference Series **299** (2011) 012014 doi:10.1088/1742-6596/299/1/012014

background with a discussion of the measured performance of the laser and then highlight a few of the user activities which leverage the performance of this unique light source. We will conclude with a discussion of the legacy and prospects for application of this development to future light sources.

2. FEL background

FELs are generally acknowledged as being re-invented by J. M. J. Madey based on the relativistic extension of earlier work of Mott and Phillips on the ubitron tube [6, 7]. The FEL utilizes a relativistic electron bunch propagating through a periodic magnetic field (wiggler or undulator) that radiates synchrotron radiation up-shifted in frequency by one Doppler shift and one Lorentz contraction of the wiggler wavelength. As a result the emitted wavelength is inversely proportional to the square of the electron beam energy (Fig. 1). Generation of 1 micron light is feasible with electron beam energies around 100 MeV. The electric field of the initial spontaneous emission combined with the wiggler field works back on the electron longitudinal distribution and causes bunching of the electrons at the optical wavelength. In this way the emitted light establishes coherence through feedback.

In an oscillator configuration mirrors are used to re-inject the previously produced light to subsequent fresh electron bunches and the intensity builds until non-linear effects take over at saturation. About 1% of the electron beam power can be extracted as optical radiation. A portion of the light is out-coupled from the optical cavity *e.g.*, through a partially transmitting mirror, and transported by mirrors to laboratories where the light may be utilized or studied. The remaining 99% of the electron beam power may be simply dumped in a water cooled copper block or, as discussed below in more detail, sent back through the accelerator 180 degrees out of the radiofrequency (rf) phase to decelerate the beam and recover the beam energy. A number of references on the physics and technology of FELs exist to which we refer the interested reader [8, 9].

Prior to the construction of the IR Demo FEL, all linac-based FELs had used pulsed machines (even the FEL on the Stanford Superconducting Accelerator - SCA because of cavity heating limitations due to electron multipacting) [10]. As a consequence they all operated at low duty factors of 10^{-3} or so [11, 12]. During the time the FEL was on, significant power was produced but low duty factor kept the true average power down and significantly limited the stability that could be achieved. Nonetheless, several US efforts [10, 11, 12] as well as several important groups outside the US [13, 14] utilized IR FEL output from RF accelerators for R&D. Although successful in delivering a number of interesting studies, independent reviews [15] remained unimpressed and FELs were relegated primarily to novelty status.

The temporal characteristics of the accelerator output dictate the temporal qualities of the laser so to achieve continuous operation of an FEL (as a continuous train of short pulses accelerated by continuous wave (CW) RF fields) it is essential to obtain a high quality CW electron source. Luckily around this time the development of Cs:GaAs photo-cathodes was deemed sufficiently mature that a DC photogun could be chosen, with temporal control obtained from a mode-locked drive laser [16]. The development of such an injector was crucial to achieving the performance of the IR Demo FEL, and its progeny, the IR/UV Upgrade. Even the present development of X-ray FELs still rests on the capability of photoinjectors initiating the process [17]. To further discuss the details of the injector system goes beyond the scope of this paper but we refer the reader to injector and gun discussions elsewhere [18].

Once the initial bunch is formed and accelerated to an energy where space charge forces no longer predominate ($\sim$10 MeV) the beam can be accelerated in standard CEBAF superconducting cavities to the energy desired for short wavelength lasing. The maintenance of electron beam quality and control of high currents is not so simple and thus our story of the accelerator system begins.

3. Accelerator developments leading to CW energy recovery

The emergence of two paradigm shifts set the stage for the evolution of FEL operation to high average power. The first addressed the FEL process itself, and was the recognition that high average powers would be most readily produced through the use of a high repetition rate drive beam rather than one with high instantaneous power. Room-temperature accelerator drivers for FELs are based on the use of extremely high single-bunch charges and high peak currents at the FEL with attendant extremely high peak laser powers. This approach is limiting because of both instabilities in the drive beam itself (due to its high instantaneous intensity) and because of beam quality limitations imposed by space charge forces and the response of the accelerator to the very large transients associated with pulsing the intense beam off and on.

The move to low peak, high average power operation was thus directly enabled by the advent of SRF technology. CW SRF accelerators for nuclear physics such as the HEPL linac, the Illinois microtron, and later, CEBAF, produce continuous streams of individual bunches, individual packets of charge, at very high repetition rate (hundreds of MHz or a few GHz) typically accelerating a bunch in every RF period. Early work in this area at HEPL [20] quickly led to the recognition that this advantage could be applied to FEL drivers. Beams with lower instantaneous currents would be of extremely high quality; so high, in fact, that the high peak current required by the FEL could be produced by beams with quite modest single bunch charge. Moreover, given the capability of SRF systems to be on all the time that is, to operate CW the average current of such a system could in fact compete with, or exceed, that available from a conventional pulsed accelerator. Success of initial CEBAF operation showed that a very high power electron drive beam of exceptional quality would be available in an SRF based system.

The second paradigm shift involved the recognition that both the efficiency and operability of a high power FEL system would be limited unless novel methods evolved for power management

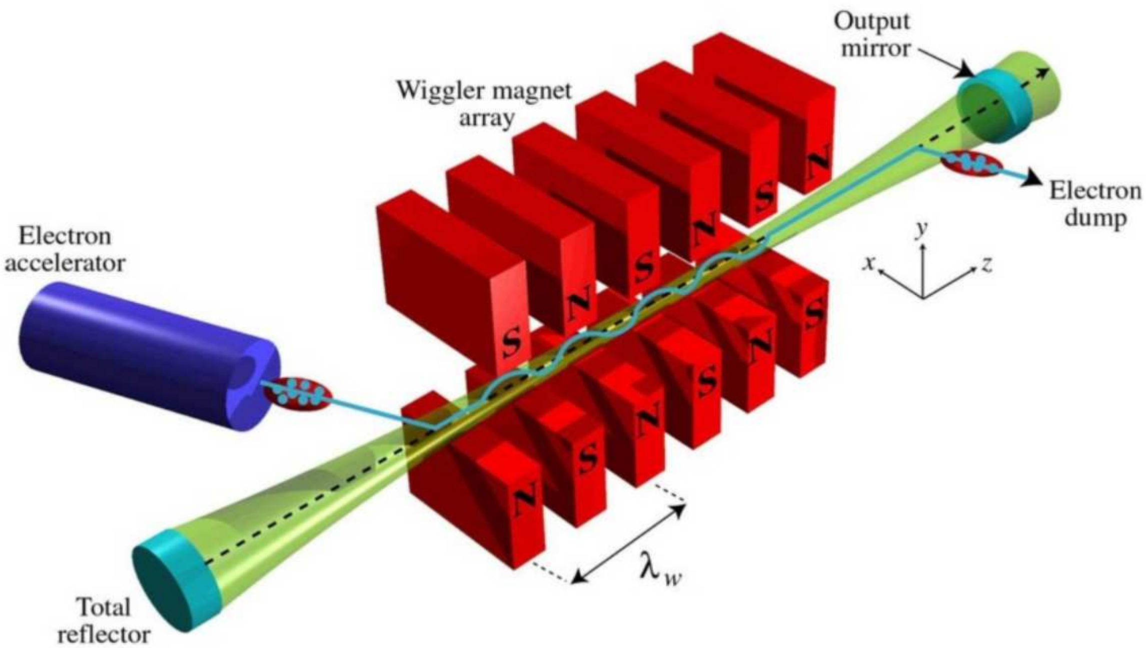

Figure 1. The FEL interaction region where bunches of relativistic electrons from the accelerator enter the wiggler simultaneously with light pulses which are amplified provided they are near the resonant wavelength.

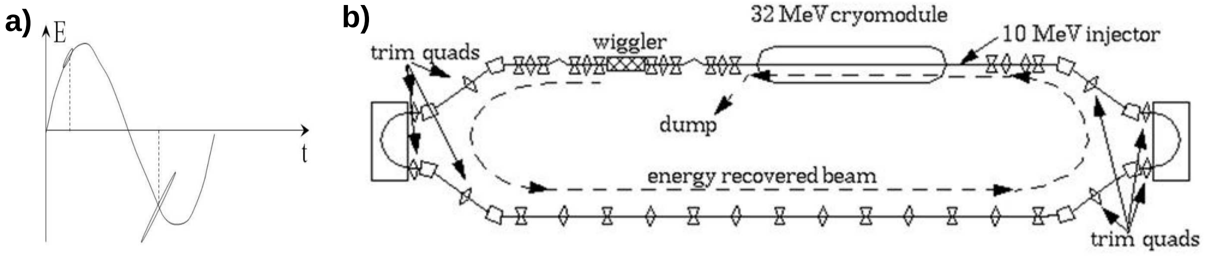

Figure 2. a) shows the timing of a beam bunch relative to the RF waveform and b) shows the accelerator geometry for energy recovery.

of the electron drive beam. FELs transform only a few percent of drive beam power into light. Kilowatt class FELs therefore require drive beams of hundreds of kilowatts of power and generate exhaust beams with similar power levels! Thus, conventional SRF linac drivers for FELs require impressive RF and wall plug power. They also demand radiation control of very high power waste beams. The expense of such a facility would be very limiting unless a viable alternative to conventional linac architectures was found. Such an enabling technology was found in energy recovery. In an effort to overcome precisely these issues, Tigner [5] had earlier proposed the notion of re-using the electron beam after it was exploited by experimentalists to put RF power back in the accelerator. Notionally, after acceleration and experimental use, the beam would be returned to the linac exactly a half-RF-period out of phase and decelerated to low energy. The power extracted from the beam would be reused for acceleration of subsequent beam. The timing of the beam with respect to the RF waveform, and a notional machine concept using this method are shown in Fig. 2. Multiple groups recognized over the next two decades the value of this technique, with an MIT group [19] performing recovery experiments with the Bates linac in the early 1980s. Even more importantly, the Stanford HEPL laboratory realizing that an SRF linac was ideally suited for this application because the absence of accelerating cavity ohmic wall losses would allow an essentially power free generation of extremely high current, high power, high quality electron beams [20].

The synergy between the success of initial CEBAF CW operation and these paradigm shifts served to further motivate efforts to pursue an FEL driven by an SRF energy recovering linac (or ERL). By the early 1990s, it had become apparent that SRF ERL drivers could be used to provide beams ideally suited for FELs and as study of this topic evolved, increasingly subtle details of a system design emerged. Three key interrelated requirements specific to CW operation quickly became apparent which have allowed the CW implementation of energy recovery and generation of extremely high FEL output powers. First of these was the emergence of a clear process of beam control through the accelerator (phase space matching) that would provide an appropriate drive beam to the FEL and recover the power from it. Secondly, the interaction of the beam with its environment and with itself must be managed to handle a menagerie of instabilities and beam-quality-degrading effects from beam currents and intensities some 25 times higher than CEBAF. Finally, essentially loss-free transport was required to handle the unprecedented beam intensities.

Solution of these demanding requirements was achieved using a novel combination of pre-existing hardware (from CEBAF) and beamline designs (largely from MIT-Bates) in a process that was truly bricolage [21]. The resulting system the IR Demo was initially operated CW in the spring of 1998 as a 50 kW single-pass linac and used in early summer 1998 to generate 155 W of CW light. Upon completion of the recirculation transport in the summer of that year, the machine was fully commissioned, acting as a true ERL, generating a drive beam of higher power than that available using only the installed RF. CW FEL performance was extended to well over 1 kW by the following summer. The follow-on system (the IR Upgrade) was a scaled-up version of the Demo; it came on line in 2003 and produced 14.3 kW CW light in the fall of 2006. It later added terahertz output ports extracting 100W in broadband femtosecond pulses and a UV FEL (UV Upgrade) producing over 100 W output. The physics considerations that went into the design of these machines are presented next.

3.1. Beam control throughout the accelerator

High power beams generated by modern CW SRF accelerators must be well controlled so as to avoid damaging levels of beam loss, typically around 1 W/m, about 10^{-6} of the total beam power. Considerable care is therefore required in the design, construction and operation of FEL driver accelerators. Beam sizes must be carefully controlled throughout the system; this requires the use of numerous focusing magnets, quadrupoles, adjusted so as to provide smooth transitions between the various regions of the machine. When properly executed and operated, control of the core beam will also provide management of halo, or beam tails, and background at large amplitude that can lead to significant beam loss.

Even more important than beam size is the control of beam energy and timing, so called longitudinal phase space management. FELs require high peak current to lase. Notionally, one might therefore try to generate and accelerate a high charge short pulse of electrons. However, such an approach would be doomed to failure by the self-fields of the electron bunch: at low energy a short, intense pulse of electrons is rapidly and severely degraded by the action of the internal repulsive forces of the charge on itself. A key advance in the design of the IR Demo was the recognition that a process of longitudinal phase space management could successfully produce, accelerate, deliver, and recover the beam needed for a high power FEL by first generating a temporally long bunch of electrons with a small energy spread. This bunch, by virtue of its length and consequentially low spatial density, would be insensitive to degradation from internal repulsive forces (the so-called space charge effects), and could be accelerated to high energy without loss of beam quality.

Acceleration occurs on the rising part of the wave form, resulting in a phase-energy correlation across the bunch, a chirp that can later be used to compress the bunch length and generate the peak current needed by the FEL. The head of the beam is at lower energy than the centroid, which is in turn at lower energy than the tail (Fig. 2a). The transport system need only be designed to bunch the chirped beam, that is, to ensure that the low energy components of the beam traverse a longer trajectory than the higher allowing the rear of the pulse to catch up with the front. The resulting short bunch can be utilized by the FEL to generate an intense optical beam; this process not only removes energy from the electrons but also smears out the beam energy spectrum resulting in a very large energy spread exhaust beam after the FEL.

A second key design advance enabling CW operation is appropriately introduced at this point of our discussion. If simply recovered, the large energy spread exhaust beam would be unmanageable when decelerated to low energy. However, the initial bunch compression process can be executed in reverse to provide energy compression during energy recovery. The large energy spread/very short bunch after the FEL would, during transport back to the linac, be decompressed, reintroducing the time-energy correlation that existed at the end of the linac after acceleration. By proper choice of transport system path length, this bunch could be re-injected exactly out of phase with the accelerated beam, resulting in recovery of beam power to the linac. Proper choice of time-energy correlation (momentum compaction) then assures that the exhaust energy, at each time, will exactly match the energy variation along the RF waveform. After recovery, at the end of the linac, the electron beam is nearly mono-energetic and can be losslessly transported into a dump. Though not novel (energy compression systems had been operated in accelerator systems prior to this time) the coupling of energy compression to the energy recovery process allows CW operation that would otherwise be precluded by beam loss

Journal of Physics: Conference Series **299** (2011) 012014 doi:10.1088/1742-6596/299/1/012014

at low energy. The process is depicted schematically in Fig. 3.

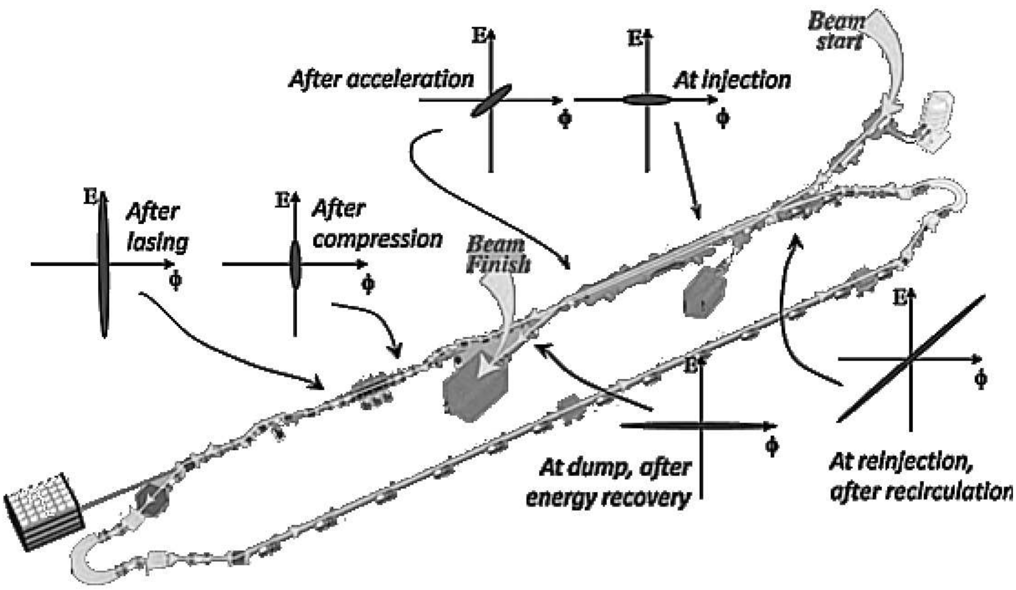

Figure 3. Schematic of longitudinal phase space management.

Extrapolation to very high powers requires use of a large energy acceptance transport system. In addition, nonlinear effects must be included such as curvature of the RF waveform and compensation for variations in bending and focusing during beam transport to the wiggler and back to energy recovery. Methods to accomplish this were confirmed for FEL-like parameters at MIT-Bates in the 1980s. The Jefferson Lab FEL linacs were designed as clones of the Bates recirculator and have repeatedly demonstrated very-low-loss operation ($< 10^{-6}$ loss). Without energy recovery, $\sim$4.5 kW of power is needed to operate each cavity while accelerating 1 mA of beam; with energy recovery, only 2 kW is needed, regardless of current.

3.2. Effects due to high beam intensity

The unprecedented CW beam intensity required by high power FEL operation stimulated concern about numerous current-driven phenomena during design and operation of the Jefferson Lab FEL drivers. A critical issue in the design of CEBAF had been the potential impact of beam break-up (BBU), a beam-intensity-driven instability that had plagued earlier generations of recirculated linacs. SRF cavities resonate at a spectrum of frequencies, not just the RF fundamental, and under certain circumstances the parasitic "higher order modes" (HOMs) can significantly influence beam behavior. In particular, if the timing of the bunch train is commensurate with the resonant frequency of an HOM, the beam can potentially put power into the mode, causing the mode to drive the beam in a manner further exciting the mode. This feedback loop can readily go unstable, causing beam loss and limiting the operating current of the accelerator.

Evaluation of this effect for the initial design indicated that it, like CEBAF, would not be susceptible to unstable behavior because of the care exercised during initial cavity design. Early experimental studies of this effect [22] stimulated a comprehensive investigation of the instability. As a result, accelerator designers are now essentially free to choose the instability threshold current [23] with well-defined engineering design criteria.

Journal of Physics: Conference Series **299** (2011) 012014 doi:10.1088/1742-6596/299/1/012014

Extrapolation to higher beam currents has required investigation and management of a number of beam-intensity-driven effects, in addition to the BBU instability. These include interactions of the bunch with itself, such as coherent synchrotron radiation (CSR), and with the accelerator environment. The latter includes the interaction of the beam electromagnetic fields (wakes) with various resonant and conductive devices in the system, resulting in localized heating. These studies point to an emergent understanding of power-flow management in high-power FEL systems. CSR-driven heating of FEL optical cavity mirrors, RF and resistive wall heating of beam-line components, and wake- and HOM-driven power deposition along the accelerator all represent potential performance limitations for the next generation of accelerator driver. Experience with the Jefferson Lab FEL drivers has shown that proper management of these effects is required for successful operation.

3.3. Management of beam loss, halo, and irregularities

Even after appropriate handling of the core beam and addressing effects related to high intensity, there are significant obstacles to high-power beam operation. Figure 4 shows a viewer image of the beam; the complexity of the spatial distribution is obvious. Historically, such structure has been attributed solely to complex nonlinear processes arising during beam formation and transport, but tests have shown the sources of halo, large amplitude beam components of intensity too low to be easily observed, but too high to be neglected, are typically quite prosaic. They include:

- drive laser light scattered over the active area of the photocathode (ameliorated by anodizing much of the active area),
- finite time response of the cathode (leading to momentum tails which are managed by the large acceptance of the transport system), and
- imperfect gating of the drive laser pulse train, leading to so-called "ghost pulses" of low intensity but well beyond the previously cited loss rate tolerance.

Experience with the Jefferson Lab drivers has shown that the halo tends to propagate independently of the core beam. Appropriate measures may thus be taken to manage it so that it is transported through the system with acceptably limited losses. Operationally, we typically measure loss patterns around the machine, and adjust focusing elements to which the core beam is insensitive but the halo responds strongly so as to alleviate excessive localized loss. We have thereby successfully achieved reliable and repeatable operation at beam powers at 1 MW levels with little loss ($< 10^{-6}$) and essentially no activation of accelerator components.

4. Lasing: IR Demo, IR Upgrade, THz, and UV output

4.1. IR Demo

Though the FEL group had great hopes for a high power FEL, one thing that cannot be stressed enough was the uncertainties in our expectations for this first machine dubbed the IR Demo. No machine like this had ever been built. The average design current for the IR Demo was 25 times that of CEBAF. The design of the FEL systems actually started with the formation of a team that worked on the design of an industrial UV laser system in 1995. The chief design principles of the IR Demo were to be as conservative and inexpensive as possible. A major issue was the uncertainty in the electron beam parameters. No one had used a photocathode DC gun with a free electron laser before or characterized injector performance at the desired charge so we had no good benchmark to use for the electron beam properties. We therefore assumed a factor-of-two safety margin over simulated values on the electron beam quality numbers, *i.e.* the transverse emittance and the bunch length were twice what the simulation codes predicted. The maximum operating energy of the accelerator would be 41 MeV based on cryomodule performance and the current would be 5 mA in the form of 135 pC bunches at 37.425 MHz. For the wiggler we chose

Journal of Physics: Conference Series **299** (2011) 012014 doi:10.1088/1742-6596/299/1/012014

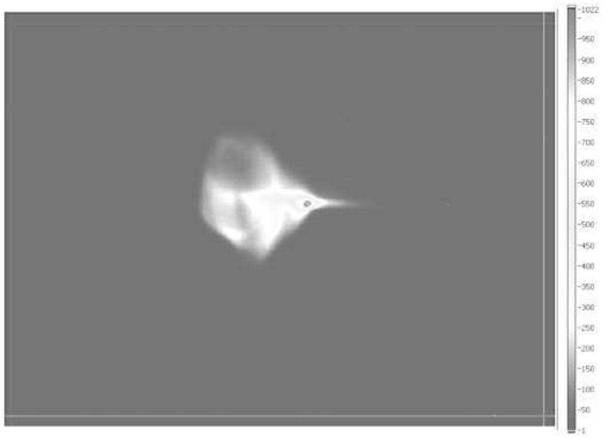

Figure 4. Beam viewer image of electron beam. The complexity of the beam structure and therefore underlying beam dynamics is clearly evident.

a proven design from the Argonne Advanced Photon Source. With a 41 MeV beam energy and a design wavelength of 3 microns we needed a fairly short wiggler period. On the other hand we were very worried about beam loss at the entrance to the wiggler so we wanted a very large gap (making achievement of the required field harder). We therefore settled on a design with a 2.8 cm period and an rms field strength of 2.7 kG. The calculated small signal gain was at least 40%, sufficient gain margin to get over 1 kW at 3.2 microns assuming perfect optical cavity mirrors. The final wiggler is shown in Fig. 5.

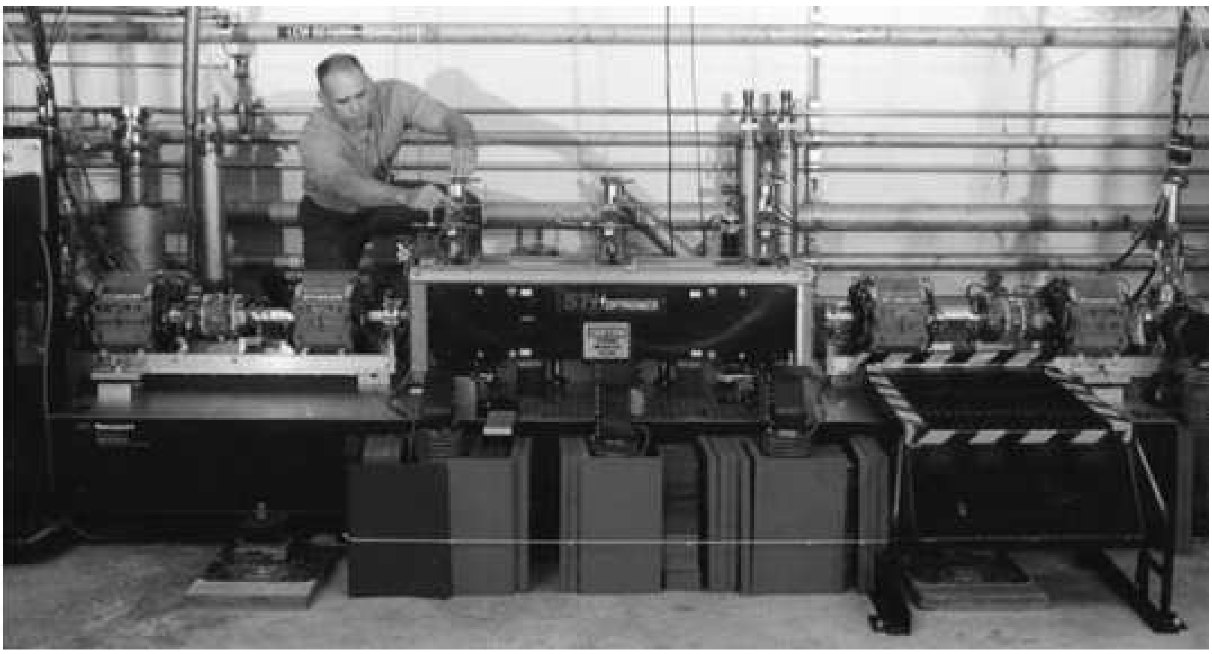

Figure 5. The IR Demo wiggler mounted on an optical bench along with four of the red matching quadrupoles. Above the wiggler one can also see three of the viewers used to align the electron beam to the optical mode.

Journal of Physics: Conference Series **299** (2011) 012014 doi:10.1088/1742-6596/299/1/012014

As the machine was commissioned we started to see payoffs from our conservative design choices. The wiggler vendor had provided us with a wiggler that could achieve the specified field with a gap much larger than the design gap. We therefore had the capability of closing the gap and using a field 50% larger than the design. Another pleasant surprise was that the beam energy from the accelerator could be as high as 48 MeV instead of 41 MeV. This meant that we could use the higher field and electron beam energy to get to the design wavelength. This increases the gain and power. To save commissioning time, we started up with only half of the design charge. This still allowed us to reach 5 mA if the repetition rate was increased to 74.85 MHz. We were surprised to find that the measured beam parameters matched the simulation values fairly closely. When all these factors were taken into account, the projected laser gain increased to about 70% per pass. With a cavity loss of 10% this provided a very comfortable gain margin for lasing.

Another essential tool in getting the laser started was the alignment system. With the help of Dick Oepts from the Felix FEL project in the Netherlands we developed a simple but very reliable method of aligning the optical cavity with both the wiggler magnetic axis and the electron beam. It is worth pointing out that the optical mode and the electron beam in the wiggler are both about the diameter of strands of spaghetti. These three-foot long strands have to be overlapped to a fraction of their diameter and one can only see them in a few locations. Nevertheless, the alignment system allowed us to routinely align the electron beam and optical cavity so that very little subsequent adjustments had to be carried out.

Because of both the alignment system and the large gain margin, the laser lased surprisingly easily when the wiggler was first installed (we had left the wiggler off the beamline until we were convinced we could actually get beam through the small gap in the wiggler vacuum chamber). Initial lasing was with 42 MeV beam at 5 microns on June 15, 1998 using our so-called “first light” optics. These only had an output coupling of 3%. We quickly ramped up to over 150 W of power in a configuration in which the beam was dumped at full energy in a straight-ahead dump (see figure 3 for a drawing of the IR Demo layout). We then changed out the output coupler for a 10% output coupler and achieved over 300 W on July 27 [24]. In straight-ahead mode we could only accelerate about 1.1 mA of beam current so the FEL power was limited to 500 W. The laser was quite stable and we could operate for long periods of time at about 30 times the previous world record power.

Once lasing had been achieved, we had to learn how to energy recover. The basic principles of this were described in the previous section. Though we knew the basic ideas we had to fill in the operational details of running this very new type of machine. We had to learn how to cope with all the problems described in the last section. Once energy recovery had been achieved, we started to ramp up the current and power and met a different limitation.

The theory of high power FEL operation with absorptive mirrors had been worked out during the design stage [25]. As the mirrors heat up their surfaces bulge out and change the waist size in the optical cavity. This distortion reduces the FEL gain and ends up clamping the power at some level. Once this power is reached it cannot be exceeded even if the electron beam current is greatly increased. The extent of the distortion for a given output power can be used to define a thermal figure-of-merit for a given mirror set. The achievable output power is proportional to the figure of merit. Our initial mirrors were made of calcium fluoride, which has a relatively poor thermal figure- of-merit. As expected, the output power with these mirrors was clamped at 500 W. Using a silicon mirror for the high reflector allowed us to reach 710 W by March 11, 1999. Finally we switched to a pair of sapphire mirrors. Sapphire has a much higher thermal figure of merit than calcium fluoride. Once we installed sapphire mirrors coated for 3 micron operation we were able to ramp the power up to 1720 W with 4.4 mA of beam current on July 15. The limitation now was just the efficiency of the laser and the available current [26].

Once we had exceeded our goal of 1000 W from the IR Demo we switched over to using the beam for user experiments. This is described in the next section. We also continued to study the machine and optimize its performance. We found that the electron beam quality was sufficient to lase at not only the third harmonic, already demonstrated on several FELs, but also at the second and fifth harmonic, which had never been accomplished [27]. We also used doubling crystals to double and triple the laser when operated at the third harmonic. This produced up to 56 W of green light and 12 W of UV light (see Fig. 6). After running the accelerator for a couple of years we were able to take advantage of the improved laser operation and energy recovery and provide 2.1 kW of laser power at 3 microns, more than a factor of two over the design power.

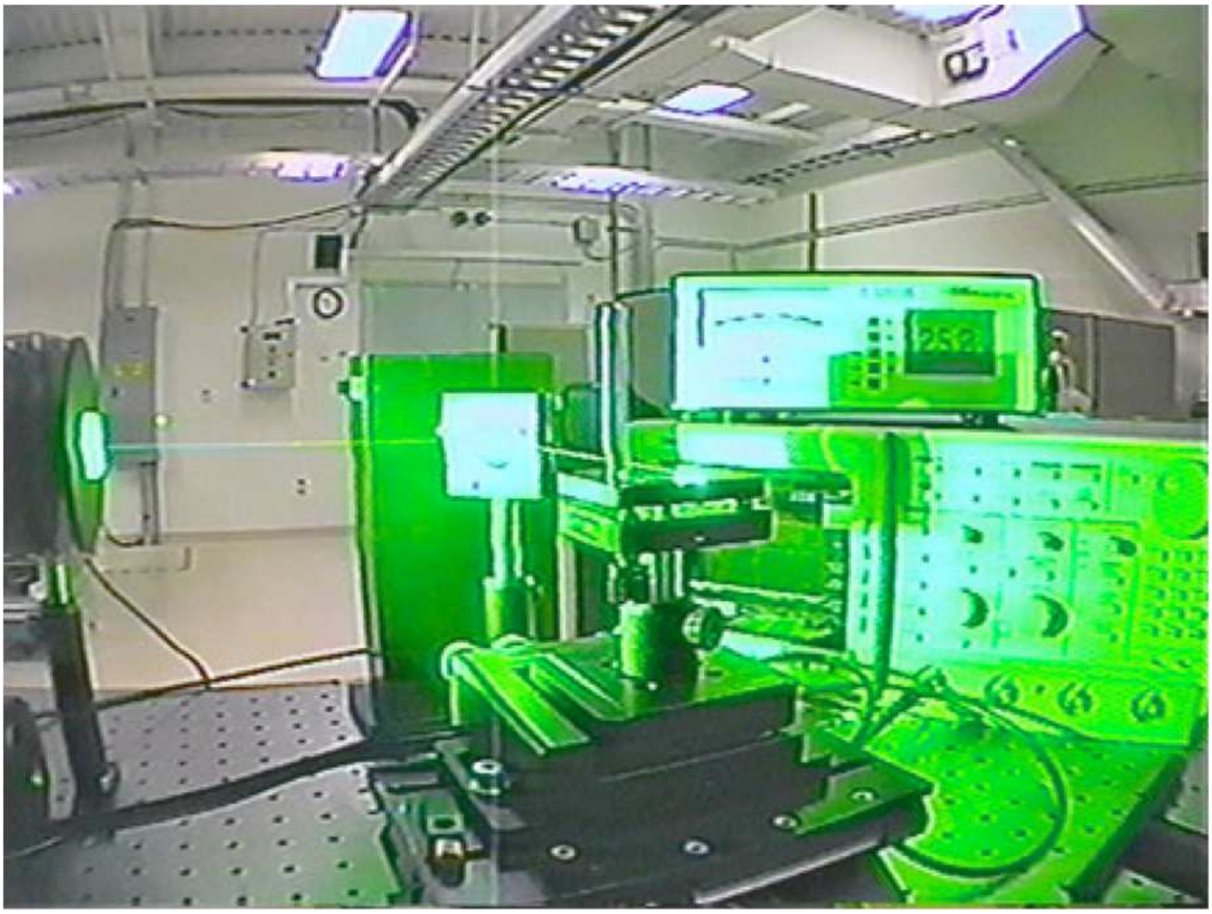

Figure 6. Frequency doubled light from the FEL. The laser was operated at 1.05 microns and doubled in a crystal to a power level up to 56 W. In this photo the 25 W beam is powerful enough to be seen going through the air.

While we delivered infrared light to users we started to look at the possibility of parasitically producing light at other wavelengths. The electron beam bunches as they enter the straight section with the FEL are very short. The synchrotron radiation produced in the last bending magnet is greatly enhanced at all wavelengths longer than the bunch length, in this case about 100 microns rms. This coherent synchrotron light is emitted in the THz part of the electromagnetic spectrum. Researchers had been working for years to produce THz radiation in a similar manner by using short pulsed lasers to produce very short electron pulses. The radiation in the 48 MeV beam was enhanced by the cube of the ratio of the electron energy to the electron rest mass energy, which is about 100 for this case. Using this simple scaling it is easy to show that we could increase the 100 μW of THz power in the laser based sources a million times to the 100 W level. To dramatically demonstrate this huge leap in power we used the THz beam to light a match [28].

The other parasitic radiation source is Thompson backscattering. The circulating light in the optical cavity collides with the electrons at the waist of the cavity where both beams are very small. The infrared photons can then scatter off the electrons and produce X-rays along the direction of the electron beam. The IR Demo produced copious amounts of these Thompson scattered X-rays and we made careful measurements of them while running the laser for other users [29].

4.2. IR Upgrade

The Navy was now interested in the next scaling step over 10 kW. If we could just triple the electron beam energy, double the electron current and keep the FEL efficiency as high as it was with the IR Demo we could get over 12 kW. This sounded like a simple, straightforward project. It turned out to be anything but. We decided at first to use an optical klystron, developed in collaboration with Advanced Energy Systems (see Fig. 7), to allow lasing in the mid-infrared with very broad tunability. This meant that we could not go to very short wavelengths unless the energy could be greatly increased. The optical klystron was also an oddity. Optical klystrons are usually used in storage rings, which have very small energy spread. An optical klystron on a storage ring has very high gain and very low efficiency. The IR Upgrade accelerator had a large energy spread and high peak current so we wanted to operate the optical klystron in a mode that had almost never been used. We used the dispersion section to just phase match between the two wigglers, creating a wiggler effectively about the same length as the actual optical klystron. This meant that the optical klystron was not the best choice since it provided less gain for the same effective number of periods compared to a simple wiggler.

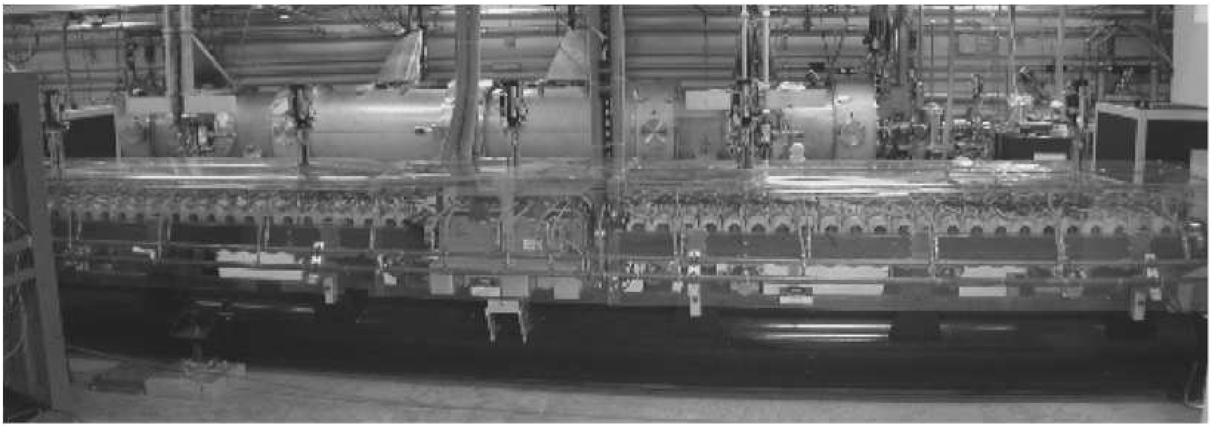

Figure 7. Optical klystron used in initial lasing with the IR Upgrade. The two wigglers in the front and back have 12 periods each. The light green dispersion section in the middle phase matches the two wigglers to make the equivalent of a 25 period wiggler. The first linac cryomodule is in the background.

Initially our energy was also not initially a factor of 3 higher than the IR Demo. The third module used to raise the energy up to over 150 MeV was not yet available. With only two cryomodules we could only operate at 80 MeV. With 10 mA of beam current that gave us 800 kW of available electron beam power. If we could achieve at least 1.25% efficiency we had a

Journal of Physics: Conference Series **299** (2011) 012014 doi:10.1088/1742-6596/299/1/012014

chance of getting to our 10 kW goal. We had achieved up to 1.5% efficiency with the IR Demo so we felt that this was an achievable goal.

One new feature we had to deal with in the IR Upgrade was bending the electron beam by 180 degrees before the FEL. This was both a blessing and a curse. It was a blessing because we could now correct for the curvature in the longitudinal phase space and produce a very short bunch. It was a curse because we had to deal with huge amounts of CSR in the final bend magnets and the resulting beam degradation. With the optical klystron operated at low energy and small dispersion, the option of operation at half charge was no longer possible. According to the FEL modeling, the Upgrade wiggler needed the full charge of 135 pC to lase effectively.

As with the IR Demo, initial laser commissioning proceeded with pulsed beam at low duty cycle. Once the cavity length had been set properly the laser lased easily. We lased at about 6 microns using zinc selenide mirrors on June 17, 2003. This system was far more flexible than the IR Demo. Up to four mirror sets could be installed at one time and the wiggler could be tuned in real-time [30]. We quickly demonstrated lasing over the full reflectivity range of the mirrors but the gain seemed much smaller than expected. This was due to the longitudinal emittance being much larger than expected. Remember that, in the IR Demo, the emittances were fairly close to simulations for a 67 pC bunch. For the 135 pC beam the transverse emittance was fairly close to the simulations but the longitudinal emittance was about four times the simulation [31]. This meant that the bunch length was much larger than we had planned. The low energy exacerbated this problem. With this poor a beam quality we could not get the 1.25% efficiency we needed.

Despite this once the second recirculation arc had been installed and CW beam was achieved we expected to ramp up quickly to many kilowatts but were frustrated to find a limit of less than a kilowatt. The problem, familiar from the IR Demo days, was mirror heating. The next 6 months were spent trying out new mirror sets to try to get mirrors with less loss. One source of loss was new, however. With the shorter bunches and higher energies possible on the IR Upgrade, the downstream mirror was bombarded with THz coherent synchrotron radiation. This was absorbed in the mirror and limited the laser power. Finally, ten months after first light we achieved 4.1 kW at 5.75 microns but we realized that we could not get to 10 kW at 80 MeV. We therefore installed the third cryomodule that provided up to 80 MeV of energy gain. This now allowed us to operate at over 150 MeV. At the higher energy the longitudinal emittance was not so much of a problem so the laser gain was quite high. Unfortunately the coherent synchrotron radiation was also enhanced so mirror heating was still a problem. We therefore reversed the optical cavity so that the CSR would land on a back-plane cooled silicon mirror instead of the zinc selenide output coupler. With this change and the larger available electron beam power we could finally get up to 8.5 kW of CW laser power. We realized that we could get even higher power if we pulsed the electron beam to avoid some mirror mount thermal issues. We obtained over 10 kW in 1 second pulses with a 1/4 Hz repetition rate on July 21, 2004. Note that we had never intended the accelerator to operate pulsed so this was quite a departure from the design operation.

Since the Navy was more interested in short wavelengths and CW operation we replaced the optical klystron first with an electromagnetic wiggler that permitted operation in the 1–2 micron wavelength range and finally a tunable permanent magnet wiggler optimized for operation in that wavelength range. The latter wiggler had both high gain and high efficiency at a wavelength of great interest to the Navy: 1.6 microns. It also allowed us to tune over a very large range with ease. With a broadband resonator using hole coupling it was possible to tune while lasing between 650 nm and 5.1 microns, more than three octaves of tuning. We could also use dielectric mirrors to easily switch between fundamental, third harmonic, and fifth harmonic lasing. The fifth harmonic lasing, achieved on the IR Demo with two high reflectance mirrors, could now be easily achieved with a 4% output coupler.

Unfortunately, at short wavelengths the mirrors are even more sensitive to mirror heating so the power at 1.6 microns was still limited to about 6 kW in CW operation. We did greatly reduce the THz heating by installing a debunching chicane and THz traps after the wiggler but there was a new problem – coherent harmonics. Most lasers produce a single wavelength that is very pure, like a flute playing a single note. The FEL naturally produces many harmonics of the fundamental lasing wavelength, similar to a violin, which produces a rich spectrum of harmonics. A prism is used to separate the 5th, 6th, and 7th, harmonics when operating at 3 microns. These harmonics allow the operator to see the beam even when the laser is operated at invisible wavelengths. When operated in the near infrared, however, significant amounts of ultraviolet radiation are produced. These UV harmonics enhanced the absorption of the fundamental light. The center of the mirror heats up by more than 10 degrees Celsius compared to the edge. The solution was to enhance the thermal figure of merit by cooling the mirrors to cryogenic temperatures. The figure of merit at liquid nitrogen temperature is at least a factor of 200 larger than at room temperature. With these mirrors the efficiency of the laser, which could routinely exceed 1.6% at low power, was now independent of power or current. We could then lase at up to 14.3 kW with 8.5 mA of current at 115 MeV.

In parallel with the IR Upgrade operations we installed an optical transport line to one of the upstairs labs for the THz radiation emitted from the last bend before the wiggler. As noted above this radiation can be well over 100 W, a unique radiation source for users. With the higher energy and current of the IR Upgrade we increased THz radiation a factor of ten more than in the IR Demo. This was used to make THz movies using THz detector arrays and cameras.

4.3. UV Upgrade

Along with the IR Upgrade we were funded to build another electron beam transport to add an ultraviolet FEL. This design benefited from all the lessons learned from the IR machine plus knowledge of the detailed machine parameter space. We used a APS undulator A prototype wiggler borrowed from Cornell that matched our design energy well. The electron beam requirements for the UVFEL include a smaller transverse emittance and a smaller energy spread. This was achieved by operating the gun at half the previous charge, 67 pC, and the highest energy available, 135 MeV. Initial operation was chosen to be at 700 nm to take advantage of relaxed tolerances and high gain. At this wavelength we achieved 165 W of output power at only 0.32 mA of average beam current. A detuning curve of more than 11 microns indicated a small signal gain in excess of 100% [32]. Since that report the lasing range has been extended down into the 438 to 362 nm range with a second mirror set. Measured gains in excess of 100% are larger than predicted by one-dimensional models and efficiencies of at least 1/2N are seen. Operation at high power is much more of a challenge than in the infrared both because the optical coating absorption is much higher and because the allowed distortion is proportional to the wavelength. Average powers in excess of order 100 W will require the use of cryogenic mirrors. Work is presently underway to fully characterize the performance, reconcile the performance with models, and utilize the harmonics of the FEL light in the VUV for photonic studies of materials.

5. FEL applications and user program

5.1. Introduction

In the early 1990s, concurrent with the technical efforts leading towards a Conceptual Design Report (CDR), advances were made to industry to drum up interest in this unique laser. It became apparent early on that there was more interest in the UV region of the spectrum than the IR, so the machine design was modified to allow future production of UV light. As it was possible to also lase in the IR, applications were also contemplated for this wavelength region. Given that the cost per photon (in units of $/kJ) for an FEL becomes comparable to more conventional laser systems once the output power is in the 50 kW range [33], applications that were initially summarized [34, 35] tended to emphasize large area processing of materials in the polymer, automotive, and electronics industries.

Why the interest in these processing applications? By the late 1980s it was clear that the use of lasers for material processing could cut costs and allow some processing steps to be done that were otherwise impossible and companies making industrial lasers begin having reasonable sales volumes of order several $100M [36]. The decade of the 1990s saw sales of all lasers used for material processing (except those based on diode lasers) quadruple, with sales exceeding $1B by the end of the decade, and continuing to grow. The majority of these lasers are CO2 or Nd:YAG technologies with CW or long-pulsed (ms long pulses at 100s of Hz rates) temporal formats, although excimer lasers and high pulse repetition frequency (PRF), Q-switched solid state lasers with ns pulsewidths have gained in popularity. Basic physics constraints on the stored energy per unit volume and gain sets the output of single-rod Nd:YAG lasers at about 1 kW, and at a few hundred watts for excimer lasers. For the processing steps performed, such as welding, cutting, and drilling, these lasers are adequate. However, for large area processing, one requires a laser with a short pulse (fs to ps) and high PRF (100 kHz to MHz) output with average power above 10 kW [37]. This is a parameter space well- matched to the specifications of an SRF ERL-based FEL. The applications can be categorized into two categories: thermal processing or ablative processing. The thermal processing can be used for surface texturing; surface amorphization; laser glazing and annealing; adhesive bond pretreatment; crystallizing amorphous silicon; laser annealing, deposition, and cutting for photovoltaics; and solvent-free cleaning. The ablative processing can be used for micromachining; cutting and slitting; and deposition of large area thin films.

For many of these processes, the benefits of using a laser were proven (and often patented) using lower powered lasers at sub-industrial scales. However, these processes haven?t become prevalent because the power of the lasers used can't be scaled to the levels ($>$ 10 kW) required. For example, despite a large number of publications showing the benefits of various short-pulsed (particularly ultrafast subpicosecond) lasers in materials processing [37, 38], their presence in a commercial environment has been limited. Ultrafast lasers have not achieved high average power status, and based on the nature of the amplification process, aren't likely to. And yet, there are compelling reasons to use ultrafast lasers; such as (1) a lower threshold for ablation, (2) more deterministic damage, (3) ablation with minimal heat-affected zone in metals, and no cracking or melting in insulators and ceramics. Along with a short-pulse time structure and the other desirable properties mentioned earlier is wavelength agility, so absorption bands (if present) in the material can be accessed.

5.2. Background

Nearly simultaneous with the start of construction of the IR Demo FEL was the appearance of two papers showing the efficacy of materials processing, specifically pulsed laser deposition (PLD) using short pulse, high pulse repetition rate (PRF) lasers [39, 40]. The idea is that the total power deposited to the surface is the same as with low (100s of Hz) laser systems, but because the energy/pulse is lower, the resulting ablation is less likely to produce droplets of material, plasma above the surface, *etc.* Experiments with carbon targets to produce diamond-like thin films, reported in [40], showed this to be the case. Nevertheless, while the use of a very short (1-20 ps) to ultrashort ($<$ 1 ps) pulse minimizes the deposition of heat to the surrounding material during the pulse, the high PRF of these lasers ensures that a fresh pulse at the same position will deposit more energy raising the temperature further before the previous pulse's energy has had time to dissipate. Work to model this effect, now called accumulation, appeared in [39] and elsewhere.

A simple spreadsheet model demonstrates this effect. Treating the laser-target interaction in the thermal equilibrium approximation, the thermal conduction problem reduces to the idealized case of a semi-infinite plane and is more amenable to a simple analytical model. This is valid for laser pulse lengths greater than about 10 ps if the optical penetration depth [given by $l_o = 2(K\tau_l)^{0.5}$, where K is the diffusivity of the material and τ_l is the pulse length] is much less than the thickness of the material. One such formulation was published by Bechtel [41] and has the solutions:

$$T(t) = \frac{2I(1-R)\sqrt{K\tau_f}}{\kappa\sqrt{\pi}}, \tag{1}$$

$$T(t) = \frac{2I(1-R)\sqrt{K\tau_f}}{\kappa\sqrt{\pi}} \left(\sqrt{\frac{t}{\tau_l}} - \sqrt{\frac{t}{\tau_l} - 1}, \right) \tag{2}$$

where I is the irradiance (in W/cm^2), R is the reflectivity of the material at the laser wavelength, and the other constants have the same meaning as above. Figure 8 shows the result of applying Eq. (1) and (2), to a target of 304 stainless steel, with a wavelength of 3 microns, a pulse repetition rate (PRF) of 20 MHz, and an irradiance of $\sim 3 \times 1011$ W/cm^2, this being a typical irradiance achieved with the IR Demo FEL.

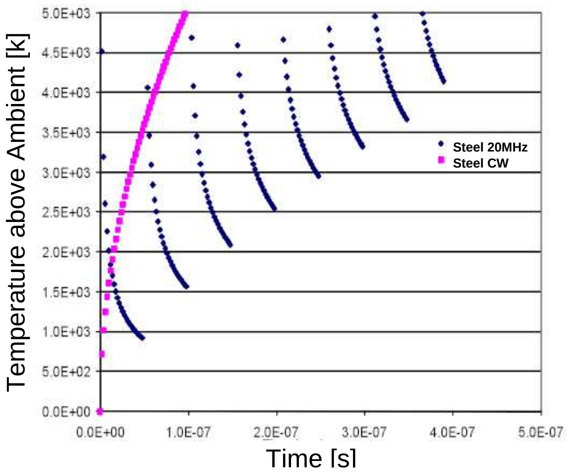

Figure 8. Modeled surface temperature of stainless steel irradiated with a high rep rate train of short pulses (blue dotted line). For comparison, the temperature profile for a CW laser with the same irradiance is shown (magenta dotted line). The parameters are given in the text.

Also shown in Fig. 8 is the surface temperature profile for CW irradiation with the same high irradiance, using the expressions published by Cohen [42]. What one sees is that, even for moderately high-reflecting metals, the first pulse takes the temperature well above the melting point. In fact, it drives the melt to vaporize, which then ionizes to form a plasma that reduces the effectiveness of subsequent pulses. Between pulses there is insufficient time for the region being irradiated to cool, so a CW-like temperature profile underlies each transient temperature excursion. This leads to the development of processing artifacts common with CW lasers, such as recast material and heat-affected zone (HAZ) on the periphery of the irradiated area. It should be noted that this was found to be true even for ablation in metals with an ultrafast laser operating at 1 kHz [43], and the same sort of techniques for mitigating the bulk heating, *e.g.*, trepanning, work equally well.

Of course, modeling is one thing, application is another. In a series of experiments, we measured the ablation rates for 304L stainless steel using the tightly focused output of the IR Demo FEL [44], this rate, as a function of plate thickness is shown in Fig. 9. The rate using an

Journal of Physics: Conference Series **299** (2011) 012014 doi:10.1088/1742-6596/299/1/012014

ultrafast Ti:sapphire laser operating at 1 kHz is also shown for comparison [45] in Fig. 9. What one sees is that for thin (< 100 micron) steel plates, the ablation rate for the Ti:sapphire laser is higher, but once the plate becomes sufficiently thick, the deeper hole traps plasma and the ablation rates become comparable. Due to the difference in PRFs (1 kHz for the tabletop laser *vs.* 18.7 MHz for the FEL), the processing rates are much higher for the FEL and are, in fact, comparable to CW lasers.

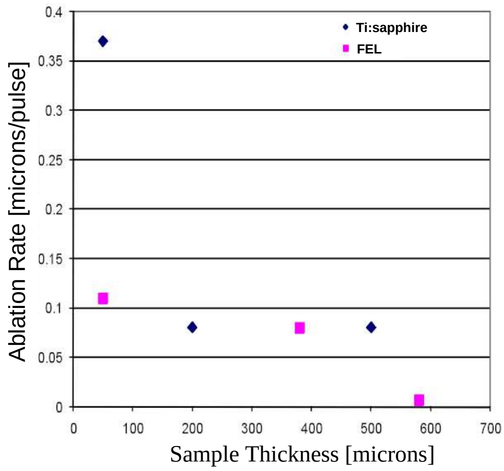

Figure 9. Comparison of ablation rates of stainless steel plates using two different ultrafast laser sources: diamonds – the FEL (wavelength = 3 microns, PRF = 18.7 MHz, F = 1 J/cm2) and squares – a Ti:sapphire laser (wavelength = 0.79 microns, PRF = 1 kHz, F = 80 J/cm2).

5.3. Applications

Carbon nanotubes are increasingly utilized in various technologies. Demand for the single-wall variety is such that the price for high-quality material is currently about \$100/gm, with pure material priced at over \$200/gm [46]. In 2003, current techniques produced about 0.2 gm/hr [47]. In comparison, worldwide demand is of order of thousands of kilograms annually [48]. Using the IR Demo FEL and delivering an average power of about 300 W onto the target, researchers obtained yields of 1.5 gm/hr, far higher than competing techniques [49]. As shown in Fig. 10, TEM (tunneling electron microscope) images show the single-wall nature of the tubes, and Raman characterization showed that the tubes were < 2 nm in diameter and had bundle sizes below 15 nm. Work to optimize yield and dimensions have continued using the IR Upgrade FEL delivering about 1 kW of 1.6 micron light into the reactor.

The processing of nonmetallic materials was also studied with the IR Demo FEL. One study [50], showed the benefits of wavelength tunability. In this investigation polyimide (DuPont Kapton HN100) was irradiated at two wavelengths, 3.1 microns (off-resonance with the amide I absorption) and at 5.8 microns (on-resonance with the amide absorption). Processing at the shorter wavelength, where the material was transparent, resulted in blackening of the material, a sign that the polymer has been thermally degraded. Tuning the laser to the longer wavelength resulted in cleanly-drilled holes, indicating the processing was nonthermal. This "cold-cutting" mode exhibited by IR FELs has been noted in earlier studies [51], enabling pulsed laser deposition (PLD) of polymers.

A comparative study of the PLD of NiFe, using either the FEL or a Ti:sapphire ultrafast laser, showed that the FEL produced superior quality films [52]. As shown in Fig. 11, the surface

Journal of Physics: Conference Series **299** (2011) 012014 doi:10.1088/1742-6596/299/1/012014

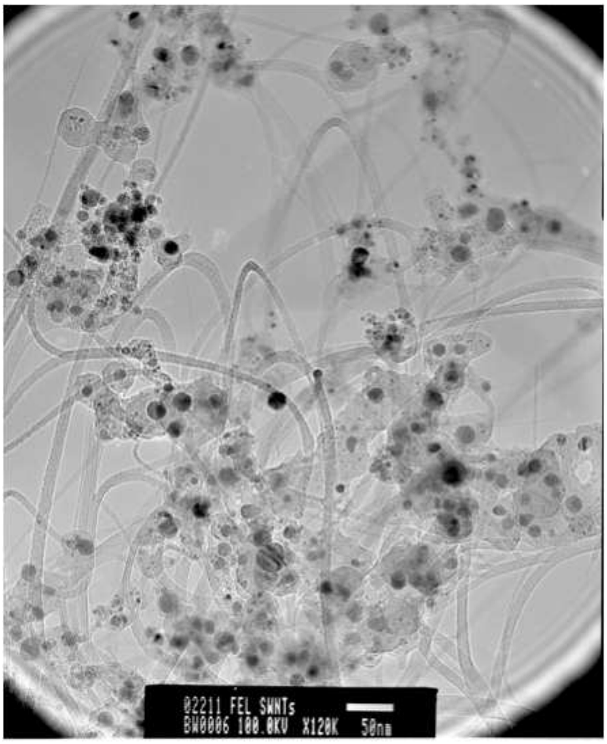

Figure 10. TEM images of single-wall carbon nanotubes produced using the IR Demo FEL operating at 3 microns.

of the films produced using the FEL is far smoother than the one made using the tabletop laser. The FEL-produced film also had a much lower coercivity. And, as one might expect from the previous discussion and data shown in Fig. 9, the deposition rate was higher for the FEL than for the tabletop laser by about 20x.

5.3.1. Surface processing Besides ablative processes, thermal (or physiochemical) processes can be employed to perform surface modifications. No material is lost, merely melted and then resolidified, or while quite hot, transformed by oxidation, nitriding, or carburization. An example which is familiar to many is the application of titanium nitride (TiN) on tool bits. This gives the bit a distinctive gold color. TiN is also used to improve the biocompatibility and wear characteristics of replacement joints [53]. In titanium nitriding studies the FEL was operated at 3 microns, usually in a burst mode (also known as a macropulse). Compared to other types of laser nitrided titanium, the FEL-produced material had a thicker and harder coating [53]. This appears to be due to the formation of oriented (200) dendrites of δ-TiN$_x$, as shown in Fig. 12. While one might think that this would result in a rough surface, in fact it is fairly smooth [54].

5.3.2. Medical applications With the wavelength and timing flexibility of the IR Upgrade FEL, it is natural to use it in medical applications. One such application is known as selective photothermolysis, the selective heating of tissues with light. By carefully choosing a wavelength in the near infrared, researchers targeted lipid-rich cells and heated them preferentially without heating the surrounding tissue. This study [55] paves the way for a laser treatment of acne, a condition where a sebaceous gland is producing too much lipid. Only the over-active glands, located a few millimeters under the surface of the skin are killed by the absorption of light, leaving

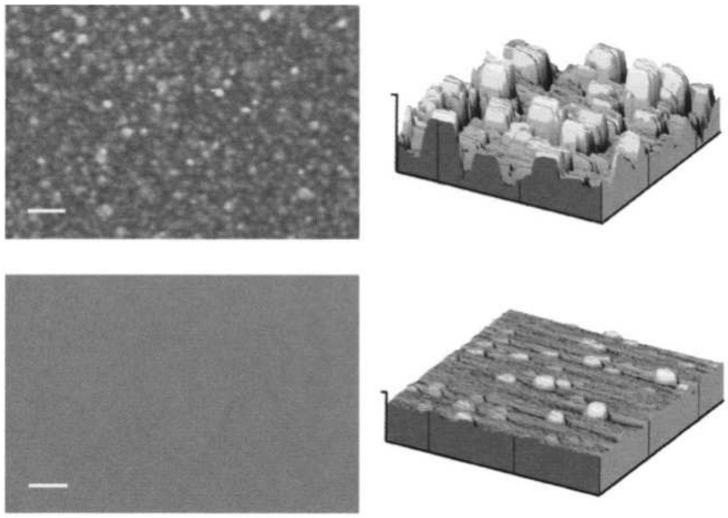

Figure 11. Scanning Electron Microscope (SEM) (left) and Atomic Force Microscope (AFM) (right) scans for thin films grown with (a) the amplified Ti:sapphire system (top pictures) and (b) the IR Upgrade FEL (bottom pictures). For the SEM scans, magnification ∼13,000x (the white line indicates 1 μmq). For the AFM scans, dimensions are 5 μm^2 and the line indicates 1500 Å height. (Figure and caption from [52].)

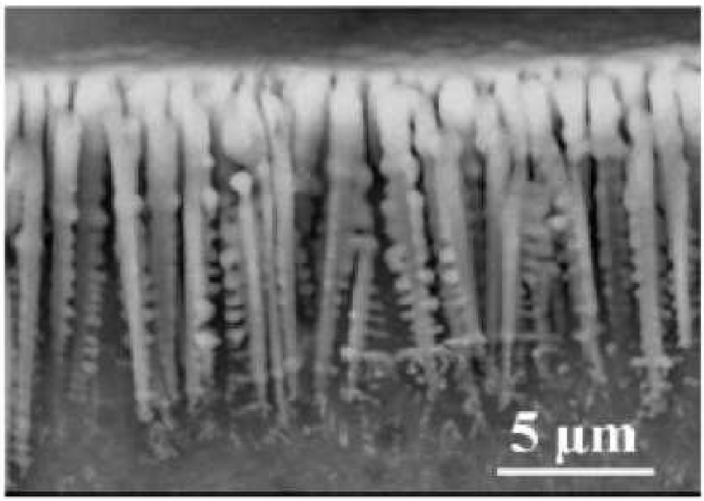

Figure 12. Formation of dendrites of δ-TiN_x after bursts of FEL irradiation at 3.15 microns [54].

the other cells unaffected. This technique also shows promise in the treatment of atherosclerosis, which if left untreated leads to heart disease and stroke.

5.3.3. Scientific applications Along with the applied research described above, the FELs ultrashort pulsewidth and wide tunability made it an excellent light source for several scientific studies. A team at the College of William and Mary used the IR Demo FEL to measure the temperature-dependent nonradiative relaxation rate of hydrogen atoms in crystalline silicon in order to determine what vibrational dynamics (*i.e.*, the coupling of the hydrogen impurity to local and lattice modes) come into play when this impurity absorbs light [56]. The high PRF, along with the relatively high average power (of order 10 W) produced data with a high SN ratio, as shown in Fig. 13.

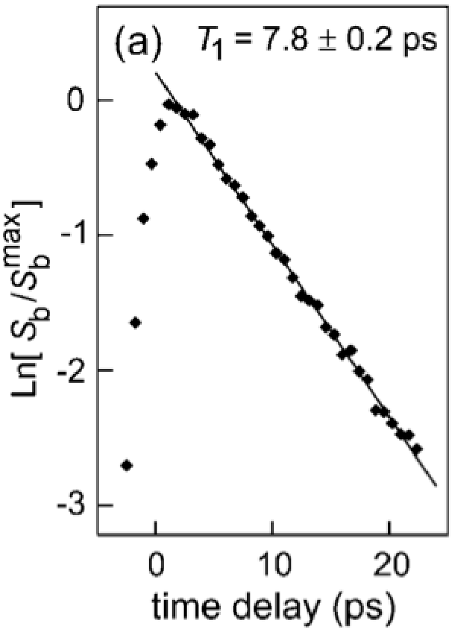

Figure 13. Decay of the transient bleaching signal from the stretch mode of the bond-centered hydrogen defect in silicon, measured at 20 K [56].

Another more recent example involves the use of the Upgrade FEL as a laser source in an effort to detect the existence of a light (sub eV) neutral boson that is a proposed dark matter candidate particle. The experiment was of the "light shining through a wall" variety, schematically diagrammed in Fig. 14. This particle, of a hypothetical family called weakly-interacting sub-eV particles (WISP), is created through an interaction between light and high magnetic field. A small fraction of photons become particles. The particles conserve momentum, so they travel in the same direction as the laser beam. Because, like dark matter, the particles couple very weakly to EM fields, they pass through the beam block and traverse another magnetic field region where they become photons once more. These are detected with an ultra low noise camera, developed for the astronomical community [57, 58, 59].

The experiment performed at Jefferson Lab was a collaborative effort, with Hampton and Yale Universities taking the lead role in scientific and results, and Jefferson Lab providing the major hardware, including the FEL. That collaboration, named the Light Pseudoscalar and Scalar Particle Search (LIPSS), took 17 hours of FEL beam, operating at a nominal wavelength of 930 nm, but detected no photons above background. Nevertheless, this result in itself set one of the most sensitive determinations of an upper bound on the scalar coupling constant (where the polarization of the FEL light is parallel to the magnetic field) for the photon-WISP conversion process [60].

5.3.4. THz applications The spectral region covered by THz radiation, roughly 1 mm to 0.1 mm in wavelength, is absorbed by transitions between electronic states and/or vibrational states in matter. Because these states are specific to a particular molecule, or ion, spectroscopy, either in

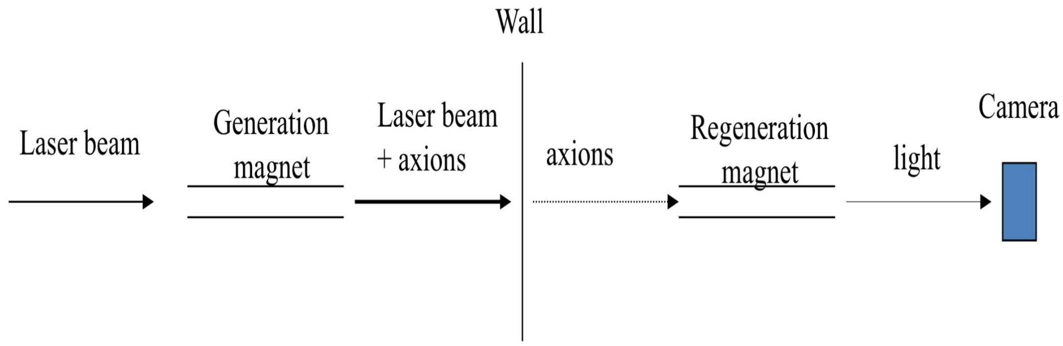

Figure 14. Schematic diagram of the dark matter candidate particle detection experiment.

the time or frequency domain can be used for location and identification. This makes the use of the THz region of the spectrum attractive for applications such as medical imaging or hazardous material detection and nondestructive testing. While the list of applications is long, progress from the lab to the field has been hampered by the lack of high average power (several watts) sources. The THz source at Jefferson Lab, is capable of producing high power, broadband, ultrashort pulsed radiation to evaluate the utility of some of these applications. It was recently used to make full-field video-frame rate images of a moving object [61]. We anticipate exciting results from the research continuing in this lab.

5.4. Conclusions

The applications discussed above are only a sampling of those pursued over the past ten years, as space does not permit us to discuss them all. The hole drilling experiments and modeling show that the FEL has a pulse format that is better suited to large area processing, whether it be via photochemical or thermal means. In addition, given the pulse format and tunablility, the FEL can be used to optimize the parameters that a tabletop laser may be configured to operate at, but don't know *a priori*. That in itself is enabling to a variety of applications that would not be readily explored, because the parameter space is too large, and the capabilities of readily available lasers are often too limited. At the time of this writing, with the addition of the UV FEL, experiments are either in the planning stages, or ready to be tested. Indeed, the prevalent view is that the best application for a FEL has yet to be found. Excitement in the field of FELs has never been higher though because of the advances into shorter and shorter wavelengths. In the past the extension of storage ring X-ray performance from first to now third generation sources has enabled an extensive set of crucial developments in material science and biology to determine static structures. However, storage rings are nearly at the apex of their possible capability. They will never be able to achieve the short pulses needed to follow chemical interactions or configuration changes of molecules. Nor will they be able to achieve the brightness required to do in vivo imaging of single proteins. The collective gain effect in FELs is expected to achieve the brightness levels needed and the application of linac rather than ring based electron sources is expected to enable the ultrashort pulses needed for dynamical studies. The developments discussed in this paper may play a major role in future light source activity by allowing high repetition rate generation of light to support many simultaneous users or achieve the sensitivity needed for examining rare events. We also hope to eventually achieve our early goal of practical use of industrial light at an affordable price for applications that cannot be achieved using conventional processing or existing lasers.

References

[1] C. Leemann, D. Douglas and G. Krafft, Ann. Rev. Nucl. Part. Sci. **51**, 413 (2001).
[2] G. A. Krafft and J. J. Bisognano, Proceedings of the 1989 IEEE Particle Accelerator Conference, 1256 (1989).
[3] R. A. Warren, Star Wars and the FEL, unpublished (1985).
[4] Laser Processing Consortium (1995) available from the FEL Dept. Office, Jefferson Lab, Newport News, Virginia 23606.
[5] M. A. Tigner, Nuovo Cim. **37**, 1228 (1965).
[6] D. A. G. Deacon, L. R. Elias, J. M. J. Madey, G. J. Ramian, H. A. Schwettman and T. I. Smith, Phys. Rev. Lett. **38**, 892 (1977).
[7] H. Motz, W. Thon and R. N. Whitehurst, Appl. Phys. **24**, 826 (1953).
[8] H. P. Freund and G. R. Neil, *Free electron generators of microwave radiation Electron Beam Generators of Microwave Radiation*, Proc. 87 Number 5 (Santa Fe, NM: IEEE) pp. 782-803 (1999).
[9] G. R. Neil and L. Merminga, Rev. Mod. Phys. **74**, 685 (2002).
[10] T. I. Smith, H. A. Schwettmann, K. W. Berryman and R. L. Swent, *Facilities at the Stanford Picosecond FEL Center*, Proc. SPIE 1854 (San Diego, CA) pp. 23-33 (1993).
[11] C. Brau, Nucl. Instrum. Meth. A **318**, 38 (1992).
[12] S. V. Benson, B. Richman and L. Vintro, Nucl. Instrum. Meth. A **296**, 110 (1990).
[13] H. Hortike *et al.*, Proceeding of the 2004 FEL Conference, THPOS17, 251 (2004).
[14] D. Oepts, R. J. Bakker, D. A. Jaroszynski, A. F. G. van der Meer and P. W. van Amersfoort, Nucl. Instrum. Meth. A **341**, 28 (1994).
[15] *Free Electron Lasers and Other Advanced Sources of Light*, National Academy Press, Washington, DC (1994).
[16] C. K. Sinclair, Nucl. Inst. and Meth. A **318**, 410 (1992).
[17] D. H. Dowell *et al.*, *Commissioning Results of the LCLS Injector*, Proc. of the 2007 FEL Conference (Novosibirsk, Russia) (2007).
[18] C. Hernandez-Garcia, P. G. O'Shea and M. L. Stutzman, Physics Today **61**, 44 (2008).
[19] J. B. Flanz and C. P. Sargent, Nucl. Inst. and Meth. A **241**, 325 (1985).
[20] T. I. Smith *et al.*, Nucl. Inst. and Meth. A **259**, 1 (1987).
[21] C. Levi-Strauss, *The Savage Mind* (Chicago, IL: University of Chicago Press) (1996).
[22] L. Merminga *et al.*, *High current effects in energy recovery linacs*, Proc. 2001 IEEE Particle Accelerator Conf. (Chicago, IL) (2001).
[23] C. Tennant C *et al.*, Phys. Rev. ST Accel. Beams **8**, 074403 (2005).
[24] S. V. Benson *et al.*, Nucl. Inst. and Meth. A **429**, 27 (1999).
[25] S. V. Benson *et al.*, Nucl. Inst. and Meth. A **407**, 401 (1998).
[26] G. R. Neil *et al.*, Phys. Rev. Let. **84** 662 (2000).
[27] G. R. Neil *et al.*, Phys. Rev. Lett. **87**, 084801 (2001).
[28] G. L. Carr *et al.*, Nature **420**, 153 (2002).
[29] J. R. Boyce, D. R. Douglas, H. Toyokawa, W. J. Brown and F. Hartemann, *Sub-picosecond, high flux, Thomson X-ray sources at Jefferson Lab's high power FEL*, Proc. 2003 Particle Accelerator Conf. (Portland, OR), 938 (2003).
[30] M. D. Shinn *et al.*, Nucl. Inst. and Meth. A **507**, 196 (2003).
[31] C. Hernandez-Garcia C *et al.*, *Longitudinal space charge effects in the Jefferson Lab IR FEL SRF Linac*, Proc. of the 2004 FEL Conference (Trieste, Italy) p. 363 (2004).
[32] R. Benson, *First lasing of the Jefferson Lab UV Demo laser*, Preprint paper MOOAI8, to appear in Proc. of the 2010 FEL Conference (Malmo, Sweden) (2010).
[33] G. R. Neil, *A cost estimation model for high power FELs*, Proc. of the 14th Particle Accelerator Conf. (Dallas, TX), 137 (1995).
[34] *High Power Ultraviolet and Infrared Free-Electron Laser for Industrial Processing Proposal* (Laser Processing Consortium) (1994) available from the FEL Dept. Office, Jefferson Lab, Newport News, Virginia 23606.
[35] *Free-Electron Lasers for Industry Vol. 1 Commercial and Technological Rationale* (Laser Processing Consortium) (1995), available from the FEL Dept. Office, Jefferson Lab, Newport News, Virginia 23606.
[36] W. M. Steen, *Laser Material Processing* 2nd ed. (New York, NY: Springer) (1998).
[37] M. J. Kelley, Proc. of SPIE **3888**, 598 (2000).
[38] M. D. Perry *et al.*, J. Appl. Phys. **85**, 6803 (1999).
[39] E. G. Gamaly, A. V. Rode and B. Luther-Davies, J. Appl. Phys. **85**, 4213 (1999).
[40] A. V. Rode, B. Luther-Davies and E. G. Gamaly, J. Appl. Phys. **85**, 4222 (1999).
[41] J. H. Bechtel, J. Appl. Phys. **46**, 1585 (1975).
[42] M. Cohen, *Material Processing Laser Handbook* (Amsterdam:North-Holland) 1578 (1972).
[43] F. Dausinger, H. Hugel and V. Konov, Proc. of SPIE **5147**, 106 (2003).
[44] M. D. Shinn, *Metal processing experiments with a high rep rate laser – the Jefferson Lab FEL*, presented at

the 5th Intl. Symp. on Laser Precision Microfabrication Conference (Williamsburg, Virginia) (2005).
[45] S. M. Klimetov *et al.*, Proc. of SPIE **4830**, 515 (2003).
[46] CheapTubes.com web site: *http://www.cheaptubesinc.com/*.
[47] M. W. Smith, private communication (2003).
[48] J. Ouellette, The Industrial Physicist **8**(6), 18 (2002).
[49] P. C. Elkund *et al.*, Nano. Lett. **2**, 561 (2002).
[50] M. J. Kelley, *Surface processing and micromachining of polyimide driven by a high average power infra-red free electron laser*, Mat. Res. Soc. Symp. (San Francisco, CA: MRS), 617 (2000).
[51] D. Bubb *et al.*, J. Vac. Sci. Tech A **19**, 2698 (2001).
[52] A. Reilly *et al.*, J. Appl. Phys. **93**, 3098 (2003).
[53] E. Carpene, P. Schaaf, M. Han, P. K. Lieb and M. D. Shinn, Applied Surface Science **186**, 195 (2002).
[54] E. Carpene, M. Shinn and P. Schaaf, Appl. Phys. A **247**, 307 (2005).
[55] R. Rox Anderson *et al.*, Lasers in Surgery and Medicine **38**, 913 (2006).
[56] M. Budde, G. Lüpke, C. Parks Cheney, N. H. Tolk and L. C. Feldman, Phys. Rev. Lett. **85**, 1452 (2000).
[57] R. Cameron *et al.*, Phys. Rev. D **47**, 3707 (1993).
[58] A. S. Chou *et al.*, Phys. Rev. Lett. **100**, 080402 (2008).
[59] C. Robilliard *et al.*, Phys. Rev. Lett. **99**, 190403 (2007).
[60] A. Afanasev *et al.*, Phys. Rev. Lett. **101**, 120401 (2008).
[61] J. M. Klopf *et al.*, Bull. Amer. Phys. Soc. **52**, X39.00013 (2007).