

**A Compact Light Source:
Design and Technical Feasibility Study of a
Laser-Electron Storage Ring X-Ray Source***

Roderick J. Loewen

Stanford Linear Accelerator Center
Stanford University
Stanford, CA 94309

SLAC-Report-632
June 2003

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

Printed in the United States of America. Available from the National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Road, Springfield, VA 22161.

* Ph.D. thesis, Stanford University, Stanford CA.

A COMPACT LIGHT SOURCE:
DESIGN AND TECHNICAL FEASIBILITY STUDY OF A
LASER-ELECTRON STORAGE RING X-RAY SOURCE

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Roderick J. Loewen

June 2003

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Prof. Ronald Ruth
(Principal Adviser)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Prof. Helmut Wiedemann
(Applied Physics and SSRL)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Prof. Todd Smith
(Physics)

Approved for the University Committee on Graduate Studies:

Abstract

Thomson scattering infrared photons off energetic electrons provides a mechanism to produce hard X-rays desirable for applied sciences research. Using a small, modest energy (25 MeV) electron storage ring together with a resonantly-driven optical storage cavity, a narrow spectrum of hard X-rays could be produced with the quality and monochromatic intensity approaching that of beamline sources at large synchrotron radiation laboratories. The general design of this X-ray source as well as its technical feasibility are presented. In particular, the requirements of optical pulse gain enhancement in an external cavity are described and experimentally demonstrated using a CW mode-locked laser.

Acknowledgments

I feel extremely fortunate to have worked at SLAC where I've had the opportunity to grow academically and professionally. I am indebted to the scientists who first taught me rf and accelerator physics—Juwen Wang, Arnold Vlioks, Samy Hanna, Sami Tantawi—as well as the SLAC faculty who supported my struggle to return to graduate school—Roger Miller and Ron Ruth. Ron was particularly brave to take in a quietly stubborn experimentalist, especially one who had interests outside the familiar terrain of high-energy physics machines, but I would not have been able to carry out these studies without Ron's commitment and keen physics insight. I am grateful for the overall support at SLAC, from the Klystron Group machine shops to the academic endorsement of the technical division leadership.

A significant amount of study on low energy electron rings was first explored by a “low emittance ring” group formed in 1999 in an attempt to find an experimental thesis based on Zhirong Huang's theoretical work on radiative damping. This group did provide the foundation for my eventual dissertation topic: Marco Venturini and Andreas Kabel studied various beam dynamics issues; Yuri Nosochkov taught me lattice design; Gordon Bowden and Jeff Rifkin provided mechanical design expertise; and Anatoly Krasnykh investigated septum magnets and fast kickers.

The success of my experimental work required the advice and help from controls experts, primarily a sympathetic Rich Abbott from LIGO (Cal Tech) with timely help from John Fox and Dmitry Teytelman at SLAC. I thank Victor Rey for looking over my shoulder to make sure I built things that worked, and asking lots of questions.

The motivation to pursue this dissertation work was inspired in part by a small cadre of GSB students—Eric Osberg, Ajay Singh, Christian Lawrence, and Erlind

Dine—who, with the guidance of Prof. Dennis Rohan, identified a possible “real world” value to this research.

Lastly, I would like to thank my wife, Jennifer, for her patience and understanding, as well as my two young sons, Benjamin and Jameson, for their unwitting ability to keep my life in perspective.

Rod Loewen
Menlo Park, CA
May 2003

Contents

Abstract	iv
Acknowledgments	v
1 Introduction	1
1.1 Background and Motivation	1
1.1.1 Synchrotron Radiation	1
1.1.2 Compact Hard X-Ray Devices	3
1.1.3 Laser Cooling of Electron Beams	4
1.2 Device Overview	5
1.2.1 Example Application: Protein Crystallography	5
1.2.2 Description of a Laser-Electron Storage Ring Source	7
1.2.3 How To Optimize X-Ray Flux	10
1.3 Focus of Dissertation	13
2 X-Rays	15
2.1 Scattering Theory	15
2.1.1 Particle View	15
2.1.2 Field View	17
2.2 X-ray Spectrum	18
2.2.1 The Laser Undulator	18
2.2.2 Electron Beam Energy Spread and Emittance	20
2.3 Laser Cooling and Quantum Excitation	20
2.4 Luminosity	22

2.4.1	Crossing Angle	23
2.4.2	Hourglass Effect	24
2.5	X-ray Window	25
3	The Electron Storage Ring and Injector	28
3.1	Electron Storage Ring System	29
3.1.1	Lattice Design	29
3.1.2	RF Cavity	33
3.1.3	Septum and Kicker System	34
3.1.4	Beam Dynamics	36
3.2	Injection System	39
4	The Optical Storage System	41
4.1	The Power Enhancement Cavity	41
4.1.1	The Fabry-Perot Cavity	42
4.1.2	Cavity Performance at Resonance	43
4.1.3	Cavity Alignment Sensitivity	44
4.1.4	Mirrors	46
4.2	Driving the Enhancement Cavity	48
4.2.1	Steering and Mode-Matching Errors	49
4.2.2	Axial Mode Matching	51
4.2.3	Dispersion Matching	56
4.3	Laser Frequency Stabilization	58
4.3.1	Characteristics of CW Mode-Locked Lasers	59
4.3.2	Frequency Stabilization Feedback Model	60
4.3.3	Pound-Drever-Hall Frequency Discriminator	62
5	Optical Cavity Experiments	66
5.1	Experimental Apparatus	68
5.1.1	Mode-Locked Laser	68
5.1.2	Transport Optics	69
5.1.3	Fabry-Perot Cavity	69

5.1.4	Frequency Stabilization Feedback	71
5.2	Experimental Procedure	73
5.2.1	Cavity Mode-Matching	73
5.2.2	Cavity Alignment	75
5.2.3	Measuring Cavity Parameters	76
5.3	Experimental Results and Discussion	80
5.3.1	Broad-band Cavity Performance	81
5.3.2	Optimizing the Feedback Servo	83
5.3.3	Narrow-band Cavity Performance	86
5.3.4	Discussion	91
	Concluding Remarks	93
	Bibliography	95

List of Tables

1.1	Target output X-ray performance.	9
1.2	Nominal values of parameters affecting luminosity.	11
3.1	Example electron storage ring parameters.	32
3.2	Magnet specifications for 25 MeV beam energy.	32
4.1	Coupling coefficients for a beam mismatched to a Fabry-Perot cavity.	51
5.1	Nominal characteristics of mirrors used for Fabry-Perot cavities.	70
5.2	Cavity mirror combinations and expected performance.	70
5.3	Measure of laser mode-matching fit to cavity.	74

List of Figures

1.1	Historical brightness of X-ray devices.	2
1.2	Concept drawing of a compact X-ray source.	8
1.3	Electron-photon beam-beam interaction.	10
2.1	Electron-photon scattering diagram.	16
2.2	Hourglass effect reduction on luminosity.	25
2.3	TiO ₂ /SiO ₂ mirror X-ray transmission characteristics.	26
3.1	Lattice functions for a 89.25 MHz repetition rate storage ring.	30
3.2	Traveling wave kicker design.	34
3.3	Lambertson septum.	36
3.4	Normalized emittance growth from IBS.	38
4.1	Basic pulse-stacking cavity.	42
4.2	Optical resonator misalignment geometry.	45
4.3	Cavity power vs. axial mode slip	53
4.4	Pulse-stacking diagram for dynamic evaluation of circulating field.	54
4.5	Dynamic peak power in a resonator for a few axial mode slips.	55
4.6	Power coupling efficiency and effective bandwidth vs. axial mode slip.	55
4.7	Cavity power vs. dispersive carrier-envelope phase slip	58
4.8	Frequency stabilization feedback system block diagram.	60
5.1	Optical table experimental layout.	67
5.2	Example of a PDH error signal	72
5.3	Early working servo transfer function.	73

5.4	Pulse ringing dynamics.	80
5.5	Broad-band cavity dip-peak measurement.	82
5.6	Broad-band cavity transmission and reflection locked signals.	83
5.7	Spectral noise of closed-loop error signal for a broad-band cavity.	83
5.8	Narrow-band cavity servo electronics sketch.	85
5.9	Optimized narrow-band cavity servo transfer function.	86
5.10	Narrow-band cavity dip-peak measurement.	87
5.11	Narrow-band cavity lock near peak axial mode.	88
5.12	Spectral noise of closed-loop error signal for a narrow-band cavity.	88
5.13	Dynamics of cavity response affecting PDH signal.	89
5.14	Peak axial mode transmission and reflection locked signals.	89
5.15	Axial mode scan using a 10.4 kHz BW cavity.	90
5.16	Fit of transmission ringing coefficients to measure cavity bandwidth.	91

Chapter 1

Introduction

1.1 Background and Motivation

1.1.1 Synchrotron Radiation

During the past three to four decades, synchrotron radiation has had a revolutionary impact on many fields of science and technology. Intense X-ray beams with wavelengths matched to the atomic scale have opened new windows to the physical and biological world. Powerful techniques such as X-ray diffraction and scattering are further enhanced by the tunability of synchrotron radiation which can exploit the subtleties of X-ray spectroscopy. One of the most dramatic applications is the detailed three-dimensional studies of protein structure using X-ray crystallography. As we enter the post-genome era, attention is focused on proteins and other macromolecules in an effort to understand disease and find new targets for drug therapies [4]. The promise of genomics is propelling structural biologists to the forefront of a large community of synchrotron light users: life science researchers are now the largest and fastest growing group of synchrotron users nationwide.

The primary source for synchrotron radiation is high energy electron storage rings, many run by large, government laboratories throughout the world. As tools of public science, these light sources are indispensable resources for pioneering research. They are also continually evolving to address ever more demanding applications. The next

generation of light sources, X-ray free electron lasers (FELs), will far exceed the existing brightness and intensities of sources today [22]. However, the increase in source quality does not necessarily translate into more access or productivity for individual researchers. Most researchers have no alternative though, since conventional X-ray source technology, like rotating anodes, have not—and likely will not—evolve at the pace necessary to compete with even the weakest beamlines. Fig. 1.1 illustrates the historical development of X-ray sources where a large void exists between conventional source technology and the dramatic improvements of insertion devices at synchrotron radiation facilities.

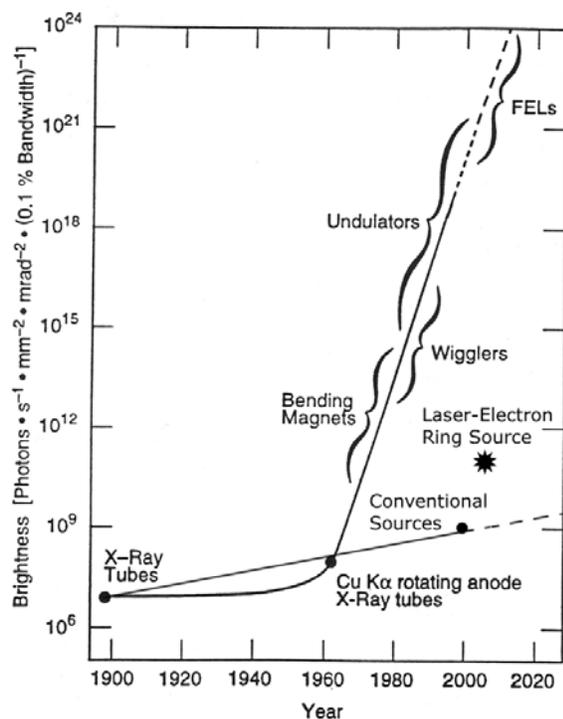


Figure 1.1: Historical brightness of X-ray devices (adapted from Winick [91]). The lack of innovation in conventional source technology has left a large gap in performance between local sources and those of large light sources. Conventional sources not only take more integrated time to perform an experiment, but they are also not tunable, and hence functionally handicapped. Even a moderate brightness synchrotron radiation source, such as the one under study, can be effectively utilized by researchers if it could be made available as a home-lab source.

Wider application of synchrotron radiation would naturally follow if compact and more affordable sources become available. For structural biologists, the centralization of research at a handful of national synchrotron radiation laboratories has led to the proposed paradigm of high-throughput protein crystallography in which a few synchrotron beamlines will be “factories” for biology research. Many life science researchers, however, would benefit from a local, on demand, synchrotron radiation

source which would be better matched to the pace and scale of existing biological laboratory research. Local sources would likely open more avenues of exploration and opportunities for experimentation to a wider set of scientists. There are also a number of medical and industrial applications that have been developed using synchrotron radiation but are impractical because no local sources yet exist with the necessary intensity and spectral properties. In particular, several groups are working on enhanced medical imaging techniques that will become clinically feasible only when a more powerful local X-ray source is available [3, 45].

1.1.2 Compact Hard X-Ray Devices

Currently only a handful of physics research ideas are candidates to develop compact and powerful X-ray sources, especially those that retain the advantageous characteristics of synchrotron radiation such as tunability and high monochromatic intensity. The compact storage rings utilized by industry for lithography have electron energies which are too low to produce X-rays needed for most material and biological applications [50]. Storage rings could be made compact, for instance with superconducting magnets, but the cost and complexity of these proposals still require more resources than most users are willing to expend [16, 93]. Of possible alternative techniques, such as channeling radiation [51] or fluorescence [52], the leading candidate for producing hard X-rays is laser-electron scattering, first investigated over a decade ago [75]. Early proposals, however, encountered difficulties, mainly from fundamental intensity limitations or designs that produce X-rays in very high-energy ranges which, again, are not suitable for the majority of current synchrotron applications.

These laser-electron researchers have been historically from one of two camps: accelerator physicists, who leverage their existing storage rings to produce high-energy X-rays or gamma rays, and laser physicists, who use advances in high peak-energy pulsed lasers. Much of the research published in the literature today comes from the latter group where table-top terawatt lasers collide with electron beams produced from pulsed linear accelerators [55, 61, 79]. These devices generate flashes of sub-picosecond X-ray pulses that have very attractive peak brightness, but are orders

of magnitude too low in average flux to make them useful to most experimenters. Some proposals mitigate average power limitations in the laser by leveraging optical storage systems (CW laser ‘supercavities’) but most still use linac technology for the electron beam source [14, 85]. Linac-based sources have two perceived advantages over a typical storage ring source: preservation of generated low emittance beams, and bunch compression in order to produce sub-ps X-ray pulses (a feat hard to do with existing light sources). The critical disadvantage, however, is average intensity.

The modest scattering cross section of colliding electrons and photons makes the design of these sources problematic. Understanding this fundamental physics limitation, however, leads to the idea of using separate storage systems for *each* particle type—a small electron storage ring and a resonant optical cavity—where the weak interaction between the two systems generates the desired X-rays. This idea was first investigated not to produce X-rays but rather to produce low emittance beams by radiative damping [44]. A follow-up study of moderate-energy electron storage rings did provide some important insights on how to optimize the ring design for hard X-ray production while identifying some critical technologies necessary to best implement the optical system.

1.1.3 Laser Cooling of Electron Beams

The photons produced by laser-electron interactions may be used to condition the electron beam through radiation damping [26, 77]. Although first proposed for high-energy GeV electron beams, one can increase the repetition rate of laser-electron scattering by using a compact, lower energy storage ring [44]. The laser cooling, offset by intra-beam scattering, results in low equilibrium emittances which could be useful for physics applications. The same storage ring, in a steady-state configuration, can be used to generate hard X-rays [11, 42]. The major advantage is high (MHz) repetition rates. The drawback with this early design is that scattered hard photons increase the equilibrium energy spread of the beam which may adversely affect beam stability as well as degrade the spectral properties of the X-rays.

Although laser damping may dominate the equilibrium beam quality, one important observation was that a high flux of X-rays could be generated without the extremely intense photon sources needed for fast damping. The storage ring could then work in a transient mode in which the injected beam retains much of its emittance quality, more similar to a proton storage ring, over many milliseconds. At these timescales, the electron bunch could then be refreshed in the ring at modest repetition rates (*e.g.* 60 Hz). This overall strategy set the stage for identifying the major feasibility issues and technological uncertainties of both the ring and the required optical storage system.

1.2 Device Overview

Before discussing the technical details of a laser-electron ring X-ray source, it is helpful to review an application that initially motivated the study of this source—protein X-ray crystallography (PX) [23]. PX is receiving a great amount of attention from public and private scientists world-wide as an integral tool for structural genomics. The following section outlines some of the general X-ray source technical requirements which are conveniently described by constraints on the X-ray phase space. This analysis leads to guidelines for designing and measuring the source's X-ray performance.

The rest of this introductory chapter highlights the conceptual design and identifies some of the major design issues which are explored in more detail in later chapters.

1.2.1 Example Application: Protein Crystallography

X-ray diffraction has been the dominant technique in solving the three-dimensional structure of macromolecules [23]. These atomic resolution studies are carried out on crystals; in principle, X-ray scattering from single molecules is possible, but the intensity is weak and the needed X-rays would destroy the sample. The crystal improves the signal-to-noise intensities of the diffraction spots which in turn reveal average positions of the non-hydrogen atoms in the protein. Only synchrotron radiation has the

intensity and spectral quality to solve most macromolecules with adequate resolution. Several thousand proteins have already had their structure solved, especially the ones easiest to put into crystal form, but there may be tens of thousands or more of interesting proteins left unsolved. The majority of proteins under study now, however, do not produce crystals that diffract well, or if they do, are physically quite small ($\sim 100\ \mu\text{m}$). The difficulties in crystallization, along with the scope of structural genomics, place special emphasis on designing future PX sources to best match the needs of structural biologists.

Technical Requirements

The quality of synchrotron radiation sources is usually described by plots of flux or brightness vs. wavelength. Whether or not higher flux or better brightness is advantageous to a particular diffraction experiment depends on the sample under test. Nave [62,63] outlines a simple measure of how to best match the source, optics, and detector requirements based on the phase space acceptance of the test specimen.

The relevant properties for a protein crystal are its dimensions, unit-cell size, and perfection. Imperfections can be described with a mosaic block model: crystal blocks with varying sizes, angular distributions, and unit-cell dimensions between (or within) the blocks. Crystals prepared through cryo-cooling techniques usually cannot retain a high degree of perfection. The practical consequence is an induced spread of the angular width of the diffracted X-ray beam (depending on the Bragg spacing). A source with much smaller angular divergence would not produce much better data than one “matched” to this angular acceptance. Nave argues that a source beam divergence $\sim 1\ \text{mrad}$ would be useful for the majority of crystals. Even larger rotation angles may be used to minimize the number of images required, for example during screening. A similar matching argument can be made with respect to source size: a smaller beam could result in local heating, or increased likelihood of radiation damage, while a larger beam would lead to increased background.

The size and angular divergence acceptance of the sample then defines a usable phase space volume. Many crystallography beamlines often quote a flux in an area \approx

200 μm square¹. Using this value for a typical crystal size, together with an angular divergence of 1 mrad, the matched transverse emittance requirement of the source is $\varepsilon = \sigma_r \sigma'_r \simeq 1 \times 10^{-7}$. The flux within this phase space volume (and desired bandwidth) is then a good measure of “useful” photon flux. For bend beamlines, or high intensity devices like wigglers, the effective emittance may be quite large; a large emittance means collimators and filters are necessary to trim the phase space by throwing away sometimes a significant amount of flux.

Lastly, there appears to be a useful range of intensity for PX experiments. Greater flux means better statistics for a set exposure time, or a shorter time for a given experiment. However, at a certain X-ray intensity, the temperature rise and gradients caused by the absorption of X-ray energy lead to radiation damage. This damage limits the effectiveness of “hot” undulator beamlines compared to more modest brightness sources proposed for PX [56]. These effects are still under study and are of topical interest in the PX community [32].

1.2.2 Description of a Laser-Electron Storage Ring Source

A conceptual picture of the X-ray source is shown in Fig. 1.2. The ring is injected with a short linear accelerator that accelerates the electron beam to the full energy desired in the ring. The electron energy necessary for 1 Å radiation is 25 MeV. The electron source produces a single electron bunch using an rf gun with a laser photocathode. The injector periodically refreshes the electron bunch in the storage ring to maintain high beam quality.

The storage ring is designed to allow the bunch to circulate in a stable fashion for about one million turns. The beam is kept tightly bunched by an rf cavity. On one side of the ring is a straight section in which the electron beam is transversely focused to a small spot. This straight section also serves as the optical gain enhancement cavity for the laser pulse. The electron bunch and the laser pulse collide each turn at the interaction point producing a burst of X-rays.

¹from the Structural Biology Synchrotron Users Organization (BioSync) compilation of beamlines in the U.S. at <http://biosync.sdsc.edu/>.

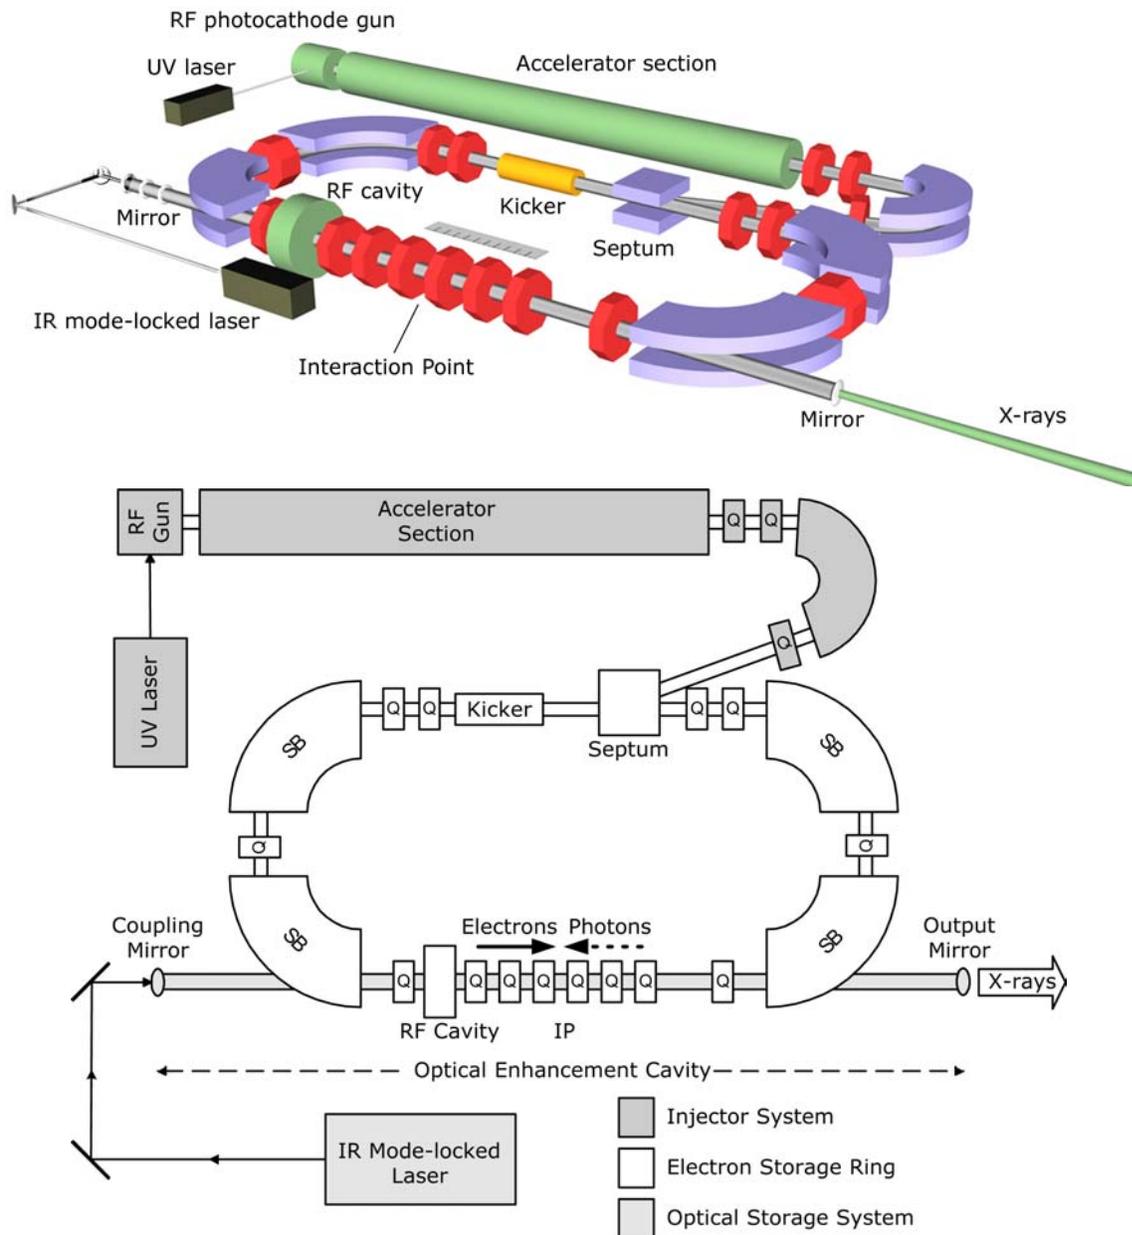


Figure 1.2: Concept drawing of a compact X-ray source. Major components are the injector (electron gun and accelerator section), the electron storage ring (shown with focusing quadrupole and bending dipole magnets), and the integrated optical cavity (between mirrors). Electron-photon scattering at the interaction point produces naturally collimated, narrow bandwidth X-rays.

Table 1.1: Target output X-ray performance.

Total Average Flux	$\gtrsim 10^{13}$ ph/s	BW $\simeq 1$
Avg. “Monochromatic” Flux*	$\approx 10^{10}$ ph/s	BW $\simeq 2 \times 10^{-4}$
Source Spot Size [†]	30 μm radius	matched beam waists
Source Divergence [†]	≈ 3 mrad	
Source Brightness	$\gtrsim 10^{11}$	ph/s/mm ² /mrad ² /0.1% BW
X-ray Energy Range [‡]	12 keV \pm 6 keV	1 \AA = 12.4 keV

*Typical value for a beamline monochromator bandwidth.

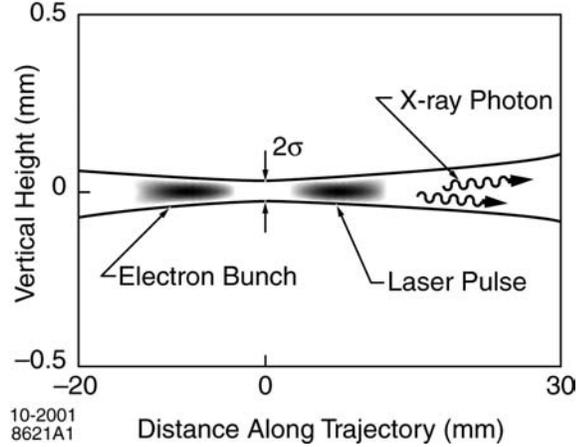
[†]Spot size and divergence correspond to $\varepsilon^n \simeq 1 \times 10^{-7}$.

[‡]Peak energy of X-rays scales as square of (tunable) electron beam energy.

A 1 μm wavelength mode-locked laser resonantly drives the enhancement cavity to build-up a high power laser pulse. The high finesse (low loss) cavity is possible due to improvements of high-reflectivity, multi-layer optics. High reliability, solid-state, mode-locked laser systems that can supply the needed laser power are also currently available.

Matched 1 cm long laser pulses and electron bunches collide to produce an X-ray spectrum equivalent to a 20,000 period undulator magnet. The X-rays are directed in a narrow cone in the direction of the electron beam as shown graphically in Fig. 1.2. They can be focused using conventional X-ray optics down to the source image size of $\sim 60 \mu\text{m}$ diameter, or slightly smaller if a larger divergence is acceptable. The target brightness and flux are shown in Table 1.1. The narrow band flux is the same magnitude as PX bend beamlines at 2nd generation storage ring sources like NSLS. Gross X-ray energy can be tuned by adjusting the electron beam energy (where fine-tuning can be accomplished by monochromator or filter adjustments). Note that the X-ray phase space follows the guidelines from the previous section: a transverse emittance of 1×10^{-7} implies a normalized transverse emittance of 5×10^{-6} for a 25 MeV beam. The configuration as described above operates with a similar photon flux up to X-ray energies of many tens of keV, and can be scaled in principle to gamma ray energies as well.

Figure 1.3: Electron-photon beam-beam interaction. The waists are shown matched, *i.e.* equivalent waist size and focus depth. For electrons, $\sigma^2 = \beta^* \varepsilon$ where $\varepsilon \approx 1 \times 10^{-7}$ and $\beta^* \approx 1$ cm determines the focus depth (and bunch length). For photons $\sigma^2 = z_R(\lambda_0/4\pi)$, so the diffraction limited emittance yields a similar focus depth $z_R \approx 1$ cm, where z_R is the Rayleigh range.



1.2.3 How To Optimize X-Ray Flux

Beam-Beam Interaction: Luminosity

Fig. 1.3 illustrates the bunch collision at the interaction point where the waists are shown matched. From an accelerator physics view, Thomson scattering can be expressed as a particle beam-beam interaction [89]. The total intensity of scattered photons can be described as a luminosity multiplied by the cross section of the event, $\dot{N}_\gamma = \mathcal{L}_0 \sigma_{\text{Th}}$. The Thomson cross section is given by

$$\sigma_{\text{Th}} = \frac{8\pi}{3} r_e^2 = 6.65 \times 10^{-29} \text{ m}^2, \quad (1.1)$$

where $r_e = e^2/4\pi\epsilon_0 m_0 c^2$ ($\simeq 2.82 \times 10^{-15}$ m) is the classical electron radius. The luminosity for two round beams with identical Gaussian distributions and waist size is

$$\mathcal{L}_0 = \frac{N_e N_L f_c}{4\pi\sigma_r^2} \quad (1.2)$$

where N_e is the number of electrons, N_L is the number of laser photons, f_c is the collision repetition rate, and σ_r is the transverse spot size. The effect of bunch lengths and crossing angles are discussed later in Sect. 2.4. Besides large numbers of particles, a high luminosity requires tight focusing and high collision rates.

The benefit of a small storage ring over linac-based sources is the very high repetition rate: ~ 100 MHz compared to ~ 100 Hz. This 10^6 increase in flux from the high

collision rate relaxes both the energy per pulse of the beams as well as the focusing requirements. The main difficulty with Thomson scattering sources—the low cross section—can in fact be used to an advantage with this kind of source. The weak interaction between the laser photons and electrons allows for a design of virtually independent storage systems: a storage ring for the electron bunches, and a resonantly driven optical enhancement cavity for the photons. In each case, the details of the storage system design, not the electron-photon interaction, are responsible for the loss mechanisms and dynamics. The efficiency for this device is dominated by the separate efficiencies of the particle storage systems rather than the scattering cross section.

Parameter Trade-offs

Parameters related to luminosity are listed in Table 1.2 along with their likely limiting factors. Of course, changing one parameter can affect others, so the task of optimizing the overall luminosity is a compromise between more flux vs. stability or power handling. A brief discussion follows for each parameter to qualitatively justify the nominal values. The design choices are examined in more detail in subsequent chapters.

Table 1.2: Nominal values of parameters affecting luminosity.

Parameter		Value	Limits
Number of electrons/bunch	N_e	1 nC	beam dynamics; injector
Laser pulse energy	U_0	1 mJ	optical cavity mirrors; input power
Focus spot size	σ_r	30 μm	stability; focusing depth
Interaction rate	f_{rep}	90 MHz	ring geometry; size
Injection rate	f_{inj}	60 Hz	emittance growth; avg. power

Number of electrons/bunch, N_e : 1 nC, or 6×10^9 electrons is the nominal charge emitted from a photo-injector cathode and is a reasonable storage ring bunch charge. Many single-bunch instabilities scale with peak current, or electron

bunch intensity, including space-charge forces. The high repetition rate also implies large average currents which can induce wakefield effects. Even without instabilities, more charge per bunch likely means faster emittance growth which dilutes brightness and may require faster re-injection rates.

Laser pulse energy, U_0 : A 1 mJ circulating pulse combined with the repetition frequency, f_{rep} , results in a circulating power ~ 100 kW. Given the nominal waist size, σ_r , and repetition rate, f_{rep} , the cavity mirrors might experience a power density approaching the maximum power density guidelines for these kinds of mirrors ($\lesssim 1$ MW/cm²). Also, under resonant conditions, the dissipated steady-state optical power is resupplied by the drive laser, which is typically available around 10 W. This implies a cavity gain enhancement of 10,000, or cavity losses of 100 ppm, which should be within technical reach.

Focus spot size, σ_r : The spot size is closely related to the desired bunch lengths through the depth of focus. A 30 μm waist corresponds $\lesssim 1$ cm long bunches. A smaller waist would imply working at shorter bunch lengths to avoid the luminosity hourglass reduction. A shorter electron bunch exacerbates beam instabilities and emittance growth effects. A shorter optical pulse also has drawbacks related to matching the optical train of the laser to the gain enhancement cavity.

Interaction rate, f_{rep} : The interaction rate is primarily set by the geometry, or optics, of the ring. A circumference of 3–4 m is a reasonable length to include an injection area, interaction point, and room for bends. Using more than one circulating pulse, or more than one interaction area, are possible options.

Injection rate, f_{inj} : Although not strictly related to luminosity, re-injection periodically refreshes the beam emittance and therefore, the quality of X-ray output. The disadvantage to a faster re-injection rate is higher rf and photocathode laser average power. Repetition rates up to 180 Hz are common for many pulsed rf sources required to power the injector system. The rate would ideally be set to some compromise of average beam quality, or emittance, vs. power consumption.

1.3 Focus of Dissertation

With the basic design parameters in place, the feasibility of a laser-electron storage ring can now be explored.

- Ch. 2 describes the X-ray output optimization and the relevant consequences to the device layout. Topics include the impact of laser radiative damping and excitation, geometrical luminosity reduction factors, calculating source brightness and spectral properties, and the feasibility of a high reflectivity IR mirror transparent to hard X-rays.
- Ch. 3 examines a sample electron storage ring design and identifies the main technological challenges. Topics include lattice design, fast injection, and some likely beam dynamics issues.
- Ch. 4 investigates the optical system from a theoretical standpoint. Topics include cavity geometry, mirror specifications, laser frequency stabilization, and resonance requirements.
- Ch. 5 experimentally demonstrates some critical aspects of the optical system. Topics include characterizing cavity performance, designing feedback stabilization circuits, and evaluating the closed-loop spectral noise of a solid-state mode-locked laser.

Although the parameters of the electron ring are unusual, designing a 25 MeV storage ring with the desired attributes should be feasible; the behavior of the beam should be easily predicted and modeled with the proven tools of accelerator physics. The purpose of the electron ring discussion is not to explore any new physics but rather to outline likely technical challenges and introduce the parameters that may impact the design of the coupled optical system.

The main technical contribution of this dissertation is the detailed study of the optical storage system: understanding how to implement a pulsed, high-gain, resonant cavity makes this X-ray source technically feasible. The experimental part of this dissertation demonstrates some critical aspects of the optical system which have not

been reported in the literature but are essential to the success of the overall design. Proving that a CW mode-locked laser is capable of resonantly driving a high-gain external cavity, much like a single-mode CW laser, leads to an effective and powerful technique to exploit laser-electron pulse interactions.

Chapter 2

X-Rays

The output of this device is a high brightness, collimated X-ray beam. The following sections describe the characteristics of the back-scattered photons, the damping effect on the beam, and ways to optimize the intensity. This analysis leads to practical considerations of how to extract the useful part of the spectrum; the X-ray phase space can then be conditioned by external optics to best match the phase space requirements of a particular application.

2.1 Scattering Theory

2.1.1 Particle View

Thomson, or Compton, scattering is often described kinematically by the classical two-particle elastic collision shown in Fig. 2.1. The relation between the scattered photon energy, E_γ , to that of the incident laser photon energy, E_L , and electron energy, E_e , is given by

$$E_\gamma = \frac{E_L(1 - \beta \cos \theta_1)}{(1 - \beta \cos \theta_2) + E_L[1 - \cos(\theta_2 - \theta_1)]/E_e}, \quad (2.1)$$

where $\beta = v/c$ is the electron velocity. For backscattered photons, $\theta_1 = \pi$ and $\theta_2 = 0$, Eq. 2.1 simplifies to the expected result

$$E_\gamma^{\text{bs}} = \frac{E_L(1 + \beta)E_e}{(1 - \beta)E_e - 2E_L} \approx \frac{E_L(1 + \beta)^2 E_e}{E_e/\gamma^2} \approx 4\gamma^2 E_L, \quad (2.2)$$

where $E_L \ll m_0c^2$, $\beta \approx 1$, and $\gamma = (1 - \beta^2)^{-1/2}$. For a fixed laser incidence θ_1 , there is a unique relation between the resultant scattered energy, E_γ , and the photon trajectory, θ_2 . Under certain circumstances, this relation can be exploited to spectrally filter the photon flux by simply collimating the scattered photons through an aperture.

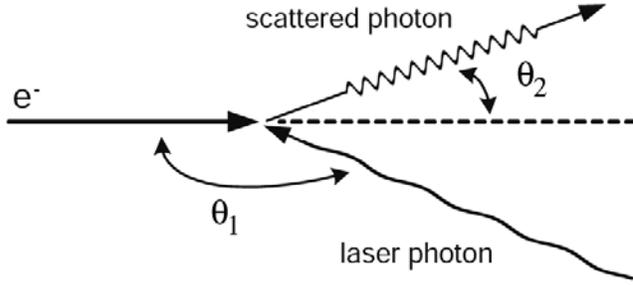


Figure 2.1: Schematic view of electron-photon scattering, with angles exaggerated. The incident angle θ_1 is close to π while the scattered photon emerges within the opening angle, $\theta_2 \lesssim 1/\gamma$.

The Klein-Nishina formula expresses the differential cross section for this scattering process [41],

$$\frac{d\sigma}{dE_\gamma} = \frac{\pi r_e^2}{2} \frac{1}{\gamma^2 E_L} \left[\frac{E_e^2}{4\gamma^2 E_L^2} \left(\frac{E_\gamma}{E_e - E_\gamma} \right)^2 - \frac{E_e}{\gamma^2 E_L} \frac{E_\gamma}{E_e - E_\gamma} + \frac{E_e - E_\gamma}{E_e} + \frac{E_e}{E_e - E_\gamma} \right], \quad (2.3)$$

where r_e is the classical electron radius, and $E_e = \gamma m_0 c^2$. This differential cross section can be used to calculate the backscattered total flux of photons N_γ in an energy bandwidth ΔE ,

$$N_\gamma \propto \frac{N_e P_0 \sigma_z}{Ac E_L} \sigma_{\Delta E} \quad (2.4)$$

where N_e is the number of electrons, P_0 is the laser pulse peak power, E_L is the photon energy, σ_z is the laser pulse length, and A is the effective interaction area [15]. The cross section $\sigma_{\Delta E}$, for instance around the peak energy $E_\gamma^{\text{max}} = 4\gamma^2 E_L$, is given

by

$$\sigma_{\Delta E} = \int_{E_{\gamma}^{\max}}^{E_{\gamma}^{\max} - \Delta E} \frac{d\sigma}{dE_{\gamma}} dE_{\gamma}. \quad (2.5)$$

2.1.2 Field View

An alternate, and perhaps more intuitive view of laser-electron scattering is to treat the counter-propagating photons as a localized electromagnetic field. An important parameter characterizing this interaction is the laser strength parameter, a_0 , which is the normalized peak amplitude of the laser's vector potential, \mathbf{A} , given by [68]

$$a_0 = \frac{e}{mc^2} \sqrt{2\langle \mathbf{A}^2 \rangle} = \frac{\lambda_0 r_e}{e} \sqrt{\frac{8\epsilon_0 I_0}{c}} \simeq 8.55 \times 10^{-10} \lambda_0 [\mu\text{m}] \sqrt{I_0 [\text{W}/\text{cm}^2]} \quad (2.6)$$

where λ_0 is the laser wavelength and I_0 is the laser intensity. For $a_0 \ll 1$, the induced transverse electron velocity (“quiver motion”) is purely linear with the laser field and can be treated classically using the non-relativistic Larmor radiation formula. For $a_0 \sim 1$, the electron motion becomes more complicated, and for higher values, non-linear. For example, a 1 mJ laser pulse, 1 cm long, focused to a 30 μm radius spot has an intensity $I_0 \approx 10^{12} \text{ W}/\text{cm}^2$; for $\lambda_0 = 1 \mu\text{m}$, this results in an $a_0 \sim 10^{-3}$ which is well within the linear region.

The laser strength parameter is analogous to the undulator (or wiggler) parameter K , usually expressed in terms of the magnet wavelength λ_u and peak field B_u as

$$K = \gamma \psi_0 = \frac{e \lambda_u B_u}{2\pi m_0 c}, \quad (2.7)$$

where K also measures the electron trajectory angle ψ_0 in units of the synchrotron radiation opening angle, $1/\gamma$ [40]. The electron beam sees an effective laser magnetic field strength

$$B_u = \frac{2}{c} \sqrt{2Z_0 I_0}, \quad (2.8)$$

where $Z_0 = 1/c\epsilon_0 = 377 \Omega$ is the free-space impedance [43]. Using the preceding example, $I_0 \approx 10^{12} \text{ W}/\text{cm}^2$ leads to an effective $B_u \approx 20 \text{ T}$. The wavelength λ_u of a counter-propagating laser pulse is simply $\lambda_0/2$. Substituting these expressions for B_u

and λ_u back in Eq. 2.7 yields Eq. 2.6, or $K = a_0 \sim 10^{-3}$. Thus the laser pulse acts just like a weak static undulator magnet, albeit one with $\sim 10^4$ periods.

2.2 X-ray Spectrum

The predicted spectral radiation power depends not only on the effective undulator parameters but heavily on the electron beam quality. In the next section, the laser “undulator” is examined first with a parallel and monochromatic beam, followed by the effects of folding in the expected beam energy spread and emittance.

2.2.1 The Laser Undulator

For an infinite-length weak undulator, there is a perfect correlation between the X-ray energy, ω , and emission angle, θ ,

$$\omega = \frac{4\gamma^2}{(1 + \gamma^2\theta^2)} \omega_0. \quad (2.9)$$

where $\omega_0 = 2\pi c/\lambda_0$ is the laser photon frequency. For small angles, a spectral filter is then equivalent to an annular ring aperture, where rings of X-ray “color” are mapped in angle away from the beam axis. The peak frequency on axis ($\theta = 0$) is defined as $\bar{\omega} = 4\gamma^2\omega_0$. The finite interaction length of the laser pulse broadens the spectral width at a fixed observation angle θ ,

$$\frac{\Delta\omega}{\omega} \simeq \frac{1}{N}, \quad (2.10)$$

where $N = 2\sigma_z/\lambda_0$ is twice the number of wavelengths over the interaction length. For a 1 cm pulse at $\lambda = 1 \mu\text{m}$, $1/N \simeq 5 \times 10^{-5}$. This spectral width is equivalent to an angular spread of $\theta = (N\gamma^2)^{-1/2} \simeq 0.1 \text{ mrad}$ near the axis, which for practical purposes can be neglected.

The radiated power per electron for a weak undulator ($K \ll 1$) is given by [40]

$$P_\gamma = \frac{4\pi\epsilon_0}{3} r_e^2 c^3 \gamma^2 B_u^2 = \frac{32\pi}{3} r_e^2 \gamma^2 I_0. \quad (2.11)$$

The angular distribution for plane undulator radiation can be expressed in terms of the this power P_γ as

$$\begin{aligned} \frac{dP}{d\Omega} &= P_\gamma \gamma^2 [F_\sigma(\phi, \theta) + F_\pi(\phi, \theta)] \\ F_\sigma(\phi, \theta) &= \frac{3}{\pi} \frac{(1 - \gamma^2 \theta^2 \cos 2\phi)^2}{(1 + \gamma^2 \theta^2)^5}, \quad F_\pi(\phi, \theta) = \frac{3}{\pi} \frac{(\gamma^2 \theta^2 \sin 2\phi)^2}{(1 + \gamma^2 \theta^2)^5}, \end{aligned} \quad (2.12)$$

where ϕ is the angle measured in the plane normal to the electron longitudinal motion (with the transverse electron motion in the $\phi = 0, \pi$ direction). Integrating this expression over the solid angle gives the fraction of total power radiated in each polarization mode as $P_\sigma = (7/8)P$ and $P_\pi = (1/8)P$. The σ -mode is strongly peaked in the forward direction $\theta = 0$ where the π -mode vanishes. A useful form of spectral power is given by integrating Eq. 2.12 over ϕ and using Eq. 2.9 [40],

$$\frac{dP}{d\omega} = \frac{3P_\gamma \omega}{\bar{\omega}^2} \left[1 - 2 \left(\frac{\omega}{\bar{\omega}} \right) + 2 \left(\frac{\omega}{\bar{\omega}} \right)^2 \right]. \quad (2.13)$$

The above relation gives the power in a given energy bandwidth accepted over all angles per electron, *i.e.* $\int_0^{\bar{\omega}} \frac{dP}{d\omega} = P_\gamma$. The angular spread for a small bandwidth, however, is naturally collimated much better than the full $1/\gamma$ opening angle. For a relative bandwidth $(\Delta\omega/\bar{\omega})$ around the peak energy, all the power falls within an angle

$$\theta_{\max} = \frac{1}{\gamma} \sqrt{\left(\frac{\Delta\omega}{\bar{\omega}} \right)}, \quad (2.14)$$

with a total integrated power

$$\int_{\bar{\omega}-\Delta\omega}^{\bar{\omega}} \frac{dP}{d\omega} \simeq 3P_\gamma \left(\frac{\Delta\omega}{\bar{\omega}} \right) - \frac{9P_\gamma}{2} \left(\frac{\Delta\omega}{\bar{\omega}} \right)^2 + O \left(\frac{\Delta\omega}{\bar{\omega}} \right)^3 \dots, \quad (2.15)$$

which shows there is a factor ≈ 3 peak in forward intensity for small bandwidths when calculating the power compared to simply multiplying the bandwidth by P_γ . For example, given a desired 1% relative bandwidth around $\bar{\omega}$ and a 25 MeV electron ($\gamma \approx 50$), the total power is $\approx 0.03P_\gamma$ and is confined within a forward cone of ≈ 2 mrad.

2.2.2 Electron Beam Energy Spread and Emittance

A real electron beam has a distribution of angular spread and energies which must be folded in to the power spectrum (Eq. 2.12). To qualitatively understand these effects, the consequence of beam energy spread and emittance can be examined separately.

First, if the beam energy spread is included, the spectral width broadens $(\Delta\omega/\bar{\omega}) = 2(\sigma_E/E_e)$. From Eq. 2.14, the natural angular spread increases as well. For the nominal design, an energy spread of 0.3% would lead to an intrinsic angular spread of ≈ 1 mrad, which is still relatively small.

Next, the angular spread can be folded in. The average angular beam divergence, θ_e is related to the normalized emittance $\varepsilon^n = \gamma\sigma_r\theta_e$. For a spot size of $30\ \mu\text{m}$, $\varepsilon^n = 5\ \mu\text{m}$, and $\gamma \simeq 50$, the natural beam divergence is $\simeq 3.3$ mrad. The beam angular spread then dominates the character of the resulting X-ray beam. The full energy spread of the X-rays, calculated from Eq. 2.9, is then $\simeq 2.7\%$.

From this analysis, the X-ray beam inherits the same emittance as the electron beam. The effective use of this spectrum will require a filter or optics designed to take advantage of this particular X-ray phase space distribution. For instance, for a fixed angular acceptance, the flux is optimized in a given energy band when the center of the band is slightly off the peak frequency $\bar{\omega}$. This source also naturally matches the requirements of large-area imaging applications [45] which need locally small divergences, *i.e.* radiation emanating from small spots. For focusing applications like protein crystallography, the $\sim 60\ \mu\text{m} \times 60\ \mu\text{m}$ spot size can be imaged or slightly magnified through appropriate X-ray optics.

2.3 Laser Cooling and Quantum Excitation

The radiative damping of the electron beam by the laser can be calculated by treating the photon emission as weak undulator radiation. If the laser pulse length is short compared to the depth of focus z_R (Eq. 4.8), P_γ can be integrated to calculate the

average energy loss per electron [43],

$$(\Delta E)_\gamma = \int P_\gamma \frac{dz}{2c} = \frac{32\pi}{3} r_e^2 \gamma^2 \frac{U_0}{z_R \lambda_0}, \quad (2.16)$$

where the laser pulse energy $U_0 = \int I_0 dx dy dz/c$. The characteristic transverse damping time is the time it takes to radiate away the initial electron energy,

$$\tau_{x,y} = \frac{E_e}{(\Delta E)_\gamma} \frac{1}{f_{\text{rev}}} \simeq \frac{1.6 \times 10^3 \lambda_0 [\mu\text{m}] z_R [\text{cm}]}{U_0 [\text{mJ}] E_e [\text{MeV}] f_{\text{rev}} [\text{MHz}]}, \quad (2.17)$$

where f_{rev} is the ring revolution frequency. Taking the nominal values of $\lambda_0 = 1 \mu\text{m}$, $z_R = 1 \text{ cm}$, $E_e = 25 \text{ MeV}$, and $U_0 f_{\text{rev}} = P_{\text{circ}} = 100 \text{ kW}$, the damping time $\tau_{x,y} = 0.64 \text{ s}$, which is much longer than the expected lifetime of the circulating bunch (*e.g.* 16 ms for 60 Hz re-injection).

The scattered photon also causes a discrete change in the electron momentum and energy. The effect of quantum fluctuations on the transverse emittance and beam energy spread have been studied in detail [26, 43]. The $1/\gamma$ opening angle of the scattered photon gives a transverse momentum recoil to the electron. By integrating over the photon spectrum and projecting onto the transverse planes, an average quantum excitation rate of the emittance due to the laser-electron interaction can be calculated as

$$\left\langle \frac{d\varepsilon_{x,y}^n}{dt} \right\rangle = \frac{3}{10} \frac{\lambda_c}{\lambda_0} \frac{(\Delta E)_\gamma}{E_e} \beta^* f_{\text{rev}}, \quad (2.18)$$

where the Compton wavelength $\lambda_c = h/m_0c \simeq 2.43 \times 10^{-12} \text{ m}$. If the beam were allowed to circulate for many damping times (and neglecting other beam dynamics), the balance between the damping rate given by Eq. 2.17 and this diffusion rate results in an equilibrium minimum emittance

$$(\varepsilon_t^n)_{\text{min}} = \frac{3}{10} \frac{\lambda_c}{\lambda_0} \beta^*. \quad (2.19)$$

For $\lambda = 1 \mu\text{m}$ and $\beta^* = 1 \text{ cm}$, this normalized emittance is $\sim 10^{-8}$, which is significantly lower than the injected normalized emittance of $5 \mu\text{m}$. In the longitudinal dimension, the relative energy change $\delta\gamma/\gamma \simeq \hbar\omega/E_e$ per photon leads to a rate of

quantum excitation to the energy spread

$$\left\langle \frac{d\sigma_E^2}{dt} \right\rangle = \frac{7}{10} \hbar \bar{\omega} (\Delta E)_\gamma f_{\text{rev}}. \quad (2.20)$$

Again, an equilibrium between this quantum excitation rate and the damping rate would result in a minimum relative energy spread

$$\left(\frac{\sigma_E}{E_e} \right)_{\text{min}} = \sqrt{\frac{7}{5} \frac{\lambda_c}{\lambda_0} \gamma} \quad (2.21)$$

which, for the nominal ring parameters, $\sigma_E/E_e \simeq 1.3\%$. Starting from an initial, injected energy spread, the growth rate to the equilibrium value is an exponential function with a characteristic time given by the damping rate. For a $\tau_{x,y} = 0.64$ s and a store time of 16 ms, the increase of the injected energy spread is $\approx 3 \times 10^{-4}$ (where the initial $\sigma_E/E_e \approx 3 \times 10^{-3}$).

Overall, for the likely working parameters of the X-ray source, the radiation damping and quantum excitation rates are much slower than the re-injection rate and should not appreciably affect the beam evolution (at least compared to other effects—see Sect. 3.1.4). If the damping time were shorter or comparable to the stored beam time, the laser interaction must be included in the beam dynamics. The main consequence would be an increased transverse emittance damping rate at the expense of a larger induced energy spread. More details can be found in Z. Huang's dissertation [42, ch. 5].

2.4 Luminosity

Luminosity is one of the most important parameters to optimize in the entire design. Since the laser intensity is too low to cause any non-linear electron motion, there is no ponderomotive or other beam-beam effects besides that of radiation damping as discussed in Sect. 2.3. The simple luminosity formula (Eq. 1.2) can then be modified by two major geometrical effects: one from a finite crossing angle, and the other from the bunch length to focal depth ratio (the hourglass effect).

2.4.1 Crossing Angle

Thomson scattering experiments are usually configured either at crossing angles of 90° to generate short (sub-picosecond) X-ray pulses, or close to 180° to optimize both the luminosity and energy transfer. The particle model predicts the energy transfer to the scattered photon as well as the relativistic $1/\gamma$ opening angle effect, but it does not address any geometrical aspects of the beams, which must be accounted for separately.

In practice, there is often some small crossing angle ϑ_c to avoid the back-scattered X-rays from impinging on an optical mirror surface. For linac-based sources, where the bunch lengths are short, this angle may not be a concern. Also, for cavities driven by CW sources, a slight crossing angle is often designed [14] or experimentally implemented [27]. For long laser pulses, however, the timing and geometrical overlap of the bunches becomes important for luminosity, which is well understood for particle colliding machines [31]. An estimate of the crossing angle sensitivity on luminosity can be evaluated by considering the angle at which the projected spot size σ_x^* in the crossing dimension becomes dominated by the angle instead of the focused waist σ_x . Using the projection,

$$\sigma_x^* = \sqrt{\sigma_x^2 + (\sigma_z \vartheta_c)^2}, \quad (2.22)$$

the angle becomes important when $\vartheta_c \sim \sigma_x/\sigma_z$. For the device parameters, $\vartheta_c = 30 \mu\text{m}/1 \text{ cm} = 3 \text{ mrad}$. If a crossing angle is chosen to avoid scattered photons hitting the laser optics, this angle should be some factor larger than the far field divergence of the laser envelope,

$$\theta_0 = \frac{\lambda_0}{2\pi\sigma_r}, \quad (2.23)$$

where $\sigma_r = w_0/2$ is the equivalent RMS spot size for the Gaussian laser waist w_0 . For a matching laser waist $\sigma_r = 30 \mu\text{m}$, the divergence $\theta_0 \simeq 5 \text{ mrad}$. If the mirrors contain several θ_0 to avoid diffraction losses in the optical cavity, the required crossing angle would dominate the projected beam size in Eq. 2.22, which further degrades the luminosity when the time structure is accounted for as well.

To avoid the crossing angle effect, some researchers propose more complicated optical systems. For instance, Tsunemi [79] built an interaction chamber in which

high-powered CO₂ laser pulses reflect off a focusing mirror with a hole for electron beam and radiation passage; the laser uses a hybrid mode pattern where the modes interfere to produce a “donut” shaped pattern at the mirror (to avoid losing photons through the central hole), yet interfere at the focal point as a nearly-Gaussian peak. Such systems are elaborate and generally unworkable for very high finesse cavities, however.

The most straightforward solution, and the one adopted for this design, is to have no crossing angle whatsoever, and let the X-rays simply pass through the optical mirror. Very high reflectivity multi-layer dielectric mirrors need $\sim 30 \lambda_0/4$ optical thickness layers, which for $\lambda = 1 \mu\text{m}$ and typical materials results in a physical thickness of $\lesssim 5 \mu\text{m}$. In theory, then, a very good reflector for IR may be a good window for hard X-rays. More about this X-ray-window/IR-mirror is discussed in Sect. 2.5.

2.4.2 Hourglass Effect

With counter-propagating bunches, the geometrical effect of the varying waists through the interaction region is also well-known [30]. The RMS bunch sizes vary from the collision point ($s = 0$) as

$$\sigma_r^2 = \sigma_r^{*2} \left(1 + \frac{s^2}{\beta^{*2}} \right), \quad (2.24)$$

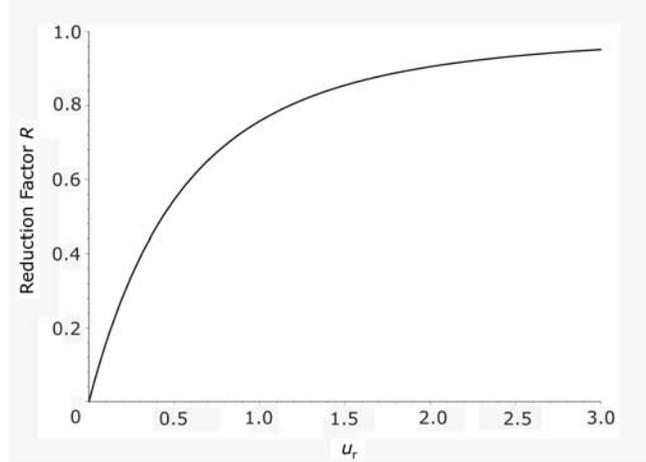
where, for symmetrical laser and electron focusing, $\beta^* = z_R$ correspond to the same transverse waist size σ_r^* at the IP. The resulting luminosity “reduction factor” is

$$R \equiv \frac{\mathcal{L}}{\mathcal{L}_0} = \int_{-\infty}^{\infty} \frac{du}{\sqrt{\pi}} \frac{e^{-u^2}}{1 + u^2/u_r^2} = \sqrt{\pi} u_r e^{u_r^2} [1 - \text{erf}(u_r)], \quad (2.25)$$

where u_r is a measure of the β -function in terms of the convoluted bunch lengths, which for the symmetric case $u_r = \beta^*/\sigma_s = z_R/\sigma_z$, and \mathcal{L}_0 is the ideal luminosity (Eq. 1.2). The function R is plotted in Fig. 2.2. The plot suggests that the bunch lengths should be comparable to or smaller than either β^* or z_R to maintain colliding efficiency. At the nominal design waists, where $\beta^* = z_R = 1 \text{ cm}$, this freedom to work at relatively large bunch lengths relaxes several otherwise detrimental effects in both systems: the smaller peak currents makes the beam dynamics in the ring more stable

(Sect. 3.1.4), the longer optical pulse makes dispersion effects in the enhancement cavity less important (Sect. 4.2.3), and the collision timing becomes less sensitive to jitter.

Figure 2.2: Hourglass effect reduction on luminosity for round beams with matching waists. The nominal working point is at $u_r = \beta^*/\sigma_s = z_R/\sigma_z = 1$ where $R \simeq 76\%$.



2.5 X-ray Window

With a 180° scattering angle, most of the usable flux will be centered on the downstream (from the electron beam) cavity mirror. From the arguments in Sect. 2.2.2, the total angle necessary to capture the majority of flux is estimated to be somewhat larger than the beam's average angular spread ($\theta_e = \varepsilon^n/\gamma\sigma_r$). For typical device parameters, this full angle should be within ~ 10 mrad, which translated ~ 1 m away at the mirror location is a ~ 1 cm diameter spot. The mirror itself will be much larger in diameter to prevent optical diffraction losses (see Sect. 4.1.4). The mirror requirement then is how to make a very high reflectivity IR mirror ($R \sim 0.9999$), yet make the central 1 cm diameter transparent to hard, but near monochromatic, X-rays.

Some general comments on high-reflectivity mirrors are discussed in Sect. 4.1.4. They are made as $\lambda_0/4$ stacks of alternating high and low index of refraction dielectrics using precision coating techniques. The low index layers are naturally less dense and contain low-Z materials, like Silica (SiO_2) or Alumina (Al_2O_3). The high index

layers often contain dense, high-Z metallic oxides, like Tantalum (Ta_2O_5) or Hafnium (HfO_2). Most of these high-Z materials have X-ray K -edge absorptions within the desired working range of hard X-ray applications (6–18 keV).

One promising combination of materials, however, is Titania/Silica ($\text{TiO}_2/\text{SiO}_2$), developed as ultra-low loss mirrors for ring-laser gyroscopes [13]. The K -edge is at 5 keV, which makes the X-ray transmission reasonable over most of the desired X-ray energy range. Fig. 2.3 illustrates the X-ray transmission losses for a typical thickness high-reflectivity stack.

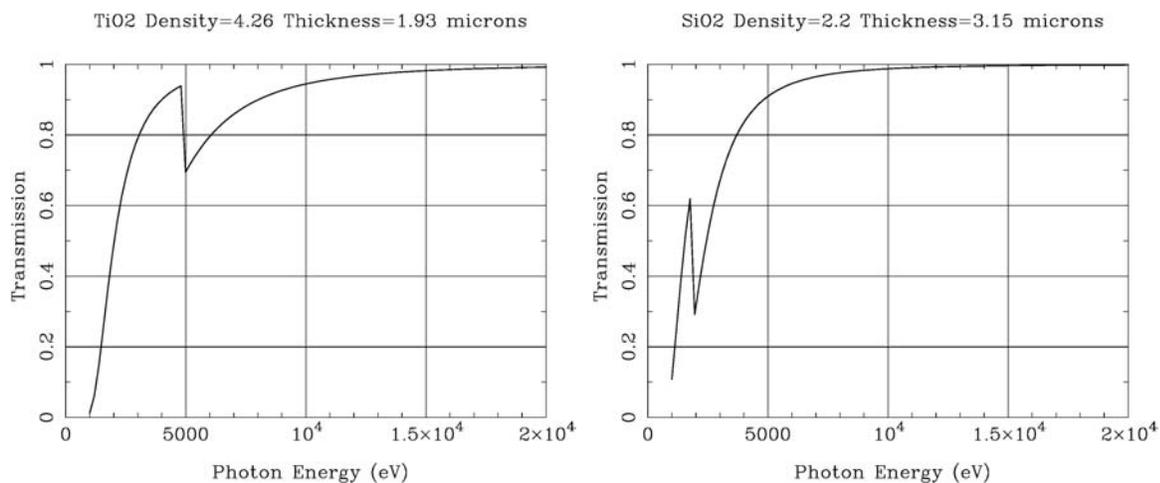


Figure 2.3: $\text{TiO}_2/\text{SiO}_2$ mirror X-ray transmission for a 32 quarter-wave layer stack. The physical thickness of each layer is $4\lambda_0/n$, where for Silica, $n_L = 1.35$, for Titania, $n_H = 2.20$, and $\lambda_0 = 1.06 \mu\text{m}$. Data and plots available from the LBL Center for X-ray Optics (www-cxro.lbl.gov/optical_constants/).

Another possible solution is to use two alternating low-Z materials that exhibit some index of refraction difference, like Alumina/Silica layers. The difficulty with such a scheme is that losses in the mirror are intrinsically higher since the optical fields must penetrate into more layers (in order to retain the high reflectivity) where power can be absorbed. The scaling for the limiting losses (for $R \sim 1$) in a quarter-wave stack is described by Koppelman's equation [5],

$$L_0 = \frac{2\pi n_0 (|k_H| + |k_L|)}{n_H^2 - n_L^2} \quad (2.26)$$

where n_0 , n_H , and n_L are the refraction indexes of the media, high-index layer, and low-index layer, while the k 's are the extinction coefficients (imaginary part of the total index of refraction). The coating technique is often more important than the material in determining the effective k value, so for a given coating technology, one wants to maximize the difference in the layer indexes. A possible hybrid solution¹ is to coat the very top layers (which contribute to the most loss) with a standard high-Z thin film and then use the low-Z materials for the remainder of the stack.

The other necessary ingredient for optical high reflectivity is to coat the multilayers on a “superpolished” substrate, like Si, SiC, or even sapphire (Al_2O_3). These examples are chosen because they are also very mechanically rigid and can be made thin over the 1 cm diameter. The usual X-ray window material, Beryllium, is typically more difficult to polish and has limited use because of regulatory safety concerns. The detailed engineering design of this mirror will likely need a close collaboration with commercial coating experts, but the overall requirements should be within technical reach. The issue of radiation damage degrading long-term optical performance is discussed in Ch. 4, Sect. 4.1.4.

¹P. Baumeister, private communication.

Chapter 3

The Electron Storage Ring and Injector

The electron storage ring and injector's function is to stably circulate a 25 MeV electron beam, transversely focus it to a waist coincident with the laser pulse at the IP, and provide a means to periodically inject \sim nC electron bunches at the full operating energy. The major ring elements consist of a magnet lattice for implementing beam optics, an rf cavity for longitudinal bunching and stability, a septum/kicker magnet system for fast single-turn on-axis injection, and beam position monitors (BPMs) or other diagnostics for beam and orbit measurements. The overall ring circumference should stay within 3–4 m to produce a high interaction rate as well as to match the typical mode-locked laser repetition frequency used to drive the optical cavity system.

Since the X-ray beam quality is mostly determined by the electron beam quality, the storage ring must be designed in a way to maintain beam emittance on timescales convenient for re-injection (\sim 60 Hz). The injector's main task is to supply a new beam pulse, at full energy, such that in a single turn the previously circulating bunch is discarded and “refreshed” by the new one. In this sense, the ring behaves more like a continuous steady-state device, although the beam dynamics of interest is in millisecond timescales, not minutes or hours as in typical high-energy storage rings.

3.1 Electron Storage Ring System

3.1.1 Lattice Design

A simple layout for the ring is to use a racetrack geometry where two straight sections offer room for injection on one side and a focusing IP on the other (Fig. 1.2 pg. 8). Since the total circumference of the ring should be kept small, the lattice is best designed with a limited number of magnets. The magnets' strengths should be specified to accommodate existing magnet designs at reasonable power consumptions (\lesssim few kW total). The beam optics must include a low- β insertion, and ideally both the IP and the injection areas are dispersion free. There is also a natural reflection symmetry folded at the IP and mid-point of the injection area which is kept for resonance stability. The more subtle criteria on the optics design include estimating and minimizing likely beam dynamics issues—primarily intra-beam scattering (IBS) and possible coherent synchrotron radiation (CSR) effects. General beam stability concerns and damping rates are discussed separately in Sect. 3.1.4.

180° Arcs The two arcs of the ring can be implemented using a double-bend achromat (DBA) design, a classic non-dispersive deflecting system [76] that only needs three elements: a pair of 90° dipole bends and one horizontal focusing quadrupole. The bend radius ρ is chosen to be as small as reasonable for obtaining a compact but power efficient dipole mechanical design.

IP The IP should focus the electron beam to a depth β^* of 1 cm by a set of quadrupole triplets, a common technique used for low- β insertions for round beams. Since the X-rays are collected far from interaction point, the quadrupoles can in principle be placed very close to the IP (longitudinally) to minimize the focusing strength and keep the beam size under control.

Matching Quads The other focusing magnets can be used as degrees of freedom to set the operating tunes ν_x and ν_y , as well as minimize the depths of other beam waists that contribute to increased intra-beam scattering and chromaticity. Sextupole

magnets, used to fix the natural negative ring chromaticity, may not be necessary for the ring's typical beam storage times (see Sect. 3.1.4).

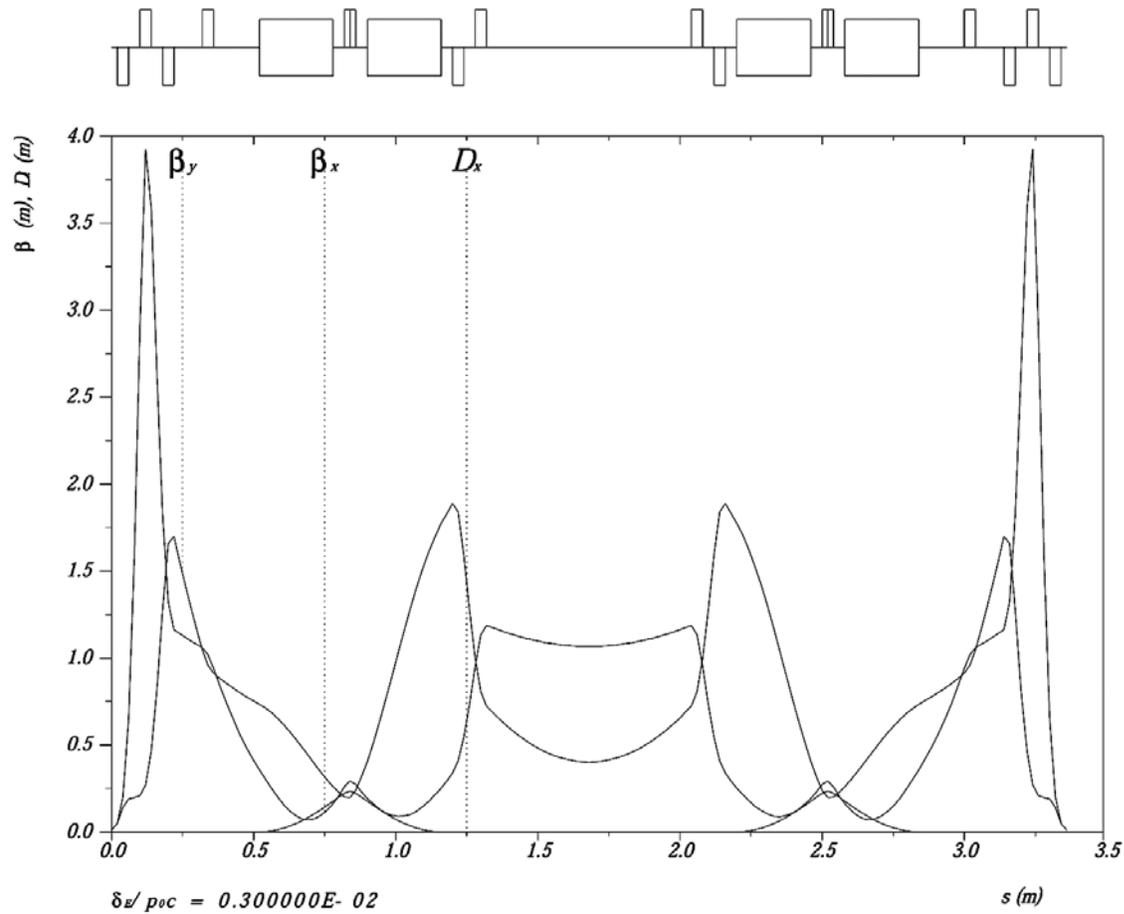


Figure 3.1: Lattice functions for a 336 cm circumference ring (89.25 MHz). Each half of the ring is symmetric from the interaction point to the middle of the injection region (center of plot). The interaction region has minimum β -functions of 1 cm in both \hat{x} and \hat{y} which matches the bunch length σ_s . There is approximately 80 cm of drift space available for the septum/kicker magnets in the injection region. The rf cavity can fit between the bend and first quadrupole magnet in the IP straight section.

Lattice Parameters

Fig. 3.1 plots the lattice functions for a 89.25 MHz revolution frequency ring design. Table 3.1 summarizes typical ring operating parameters. The design assumes an rf system running at L-band, 1428 MHz, to keep the bunch length σ_s matched near the value of the focal depth $\beta^* = 1$ cm (see Sect. 3.1.2). The operating beam is round at the IP, has an injected normalized transverse beam emittance of $\lesssim 5 \mu\text{m}$, and a relative energy spread of 0.3%. The radiation damping time, including the contribution from the ~ 1 mJ circulating laser pulse, is still an order-of-magnitude longer than the expected re-injection time. As such, the stored beam retains much of the properties of the injected beam rather than reaching the usual equilibrium values in typical radiation-dominated electron storage rings.

Magnets

All magnets can be electromagnets in order to scale the fields commensurate with the operating range of electron energies. A summary of magnet specifications are listed in Table 3.2. The expected designs stay well within recommended maximum pole-tip fields of 8–10 kG [89]. The dipole magnets used in the lattice (Fig. 3.1) are combined function magnets, although with relatively modest focusing strength. The short quadrupole magnets will have fringing fields but the aberrations are in fact well understood and smaller in magnitude than found in many conventional small aperture quadrupoles [86].

The beam vacuum chamber could be 2 cm diameter stainless steel tube. The vertical chamber height through the bends may be rectangular, however, to reduce the vertical height h for better CSR shielding (see Sect. 3.1.4) while keeping the cross-sectional area large for vacuum pumping. The overall beam conductance in the system would likely be adequate to maintain $\lesssim 10^{-8}$ torr scale vacuum, which, from a beam lifetime perspective, is more than adequate.

Table 3.1: Example electron storage ring parameters.

Parameter		Value
Electron Energy	E_e	25 MeV
Number of Electrons	N_e	6×10^9 (1 nC)
Revolution Frequency	f_{rev}	89.25 MHz
Average Beam Current	I_{avg}	86 mA
Normalized Transverse Emittances	$\varepsilon_x^n, \varepsilon_y^n$	5 mm mrad
Betatron Tunes	ν_x, ν_y	2.33, 1.39
Chromaticity	ξ_x, ξ_y	-8.9, -3.8
Momentum Compaction Factor	α	0.12
Synchrotron Energy Loss/Turn	$(\Delta E)_0$	0.21 eV
Laser ($U_0 = 1$ mJ) Energy Loss/Turn	$(\Delta E)_\gamma$	0.38 eV
Total Transverse Damping Time	$\tau_{x,y}$	0.47 s
RF Cavity Frequency	f_{cav}	1428 MHz
Harmonic Number	h	16
RF Cavity Voltage	V_{cav}	30 kV
RF Momentum Acceptance	$(\Delta p/p_0)_{\text{max}}$	2 %
Synchrotron Tune	ν_s	0.018
RMS Injected Energy Spread	σ_E/E_e	0.3 %
RMS Bunch Length	σ_s	1 cm

Table 3.2: Magnet specifications for 25 MeV beam energy.

Dipole Magnets (4 ea.)		Quadrupole Magnets (14 ea.)	
Length	26.0 cm	Effective length	4.0 cm
Bend radius ρ	16.6 cm	Minimum spacing	4.0 cm
\mathbf{B} field	0.5 T	Maximum $ k $	676 m^{-2}
Focusing strength k	-11.4 m^{-2}	Maximum g	56 T/m
Power consumption	≤ 1 kW		

3.1.2 RF Cavity

The rf cavity supplies the longitudinal restoring force in order to maintain the electron bunch length. The choice of frequency and operating voltage depends on the relations of bunch length, momentum acceptance, and geometrical size. A low rf frequency helps create a large acceptance (bucket height) for a given voltage while naturally producing long bunch lengths. From a size perspective, a pill-box π -mode cavity would take up $\sim \lambda/2$ in length, so there is a practical limit on how low a frequency to use. For example, a frequency of 1428 MHz has a $\lambda/2$ length of ~ 10 cm. The actual cavity space required in the ring circumference may be less than this value depending on the particular cavity design, which can trade-off gap spacing, cavity Q, and drive power. Shorter, loaded (non-resonant) cavities are also possible at even much lower frequencies. At the expected few kW power levels, CW sources at L-band are commercially available. Typical operating parameters are included in Table 3.1.

The design of rf cavities in storage rings is well understood, the details of which are listed in standard texts [89,90]. The particular cavity used in this ring has the unusual feature of being heavily beam-loaded: the induced voltage by the beam is much greater than the net required cavity voltage for stable operation. However, since the effective cavity voltage is the phasor sum of both the beam induced voltage and the generator voltage, there are standard methods to adjust the relative tuning of the rf to cavity frequencies in order to insure proper stable phase. The practical consideration for large beam loading is that there must be a way to handle any transients. Since the cavity decay time is large compared to the ring revolution time, the only transients occur at start-up (no beam). In this case, the transient can be dynamically compensated using a fast phase shifter on the rf, or the beam intensity and cavity voltage can be ramped up over several injection cycles until steady-state powers are reached.¹

¹A. Hofmann, private communication; phase compensation modeled by A. Kabel (SLAC)

3.1.3 Septum and Kicker System

For on-axis, single-turn injection, the most stringent requirement on the injection magnet system is that the fields generated to kick in a new bunch are gone by the time that bunch comes back around in one circulation time (~ 10 ns). One candidate kicker is a traveling-wave design where an electromagnetic wave travels opposite to the electron beam on a transmission line formed by two parallel plates (Fig 3.2). The minimum fill-time (and fall-time) is defined as $\tau = 2l/c$ where l is the length of the electrodes. Feed-thru ceramic connectors are used for input and output to loads which are terminated at the characteristic line impedance (*e.g.* 50Ω).

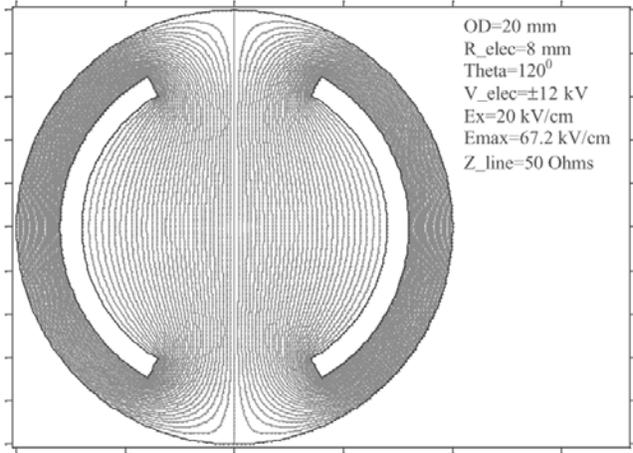


Figure 3.2: Traveling wave kicker design for a 2 cm diameter beam pipe. For 30 cm long electrodes, the rise and fall times are ~ 2 ns and the integrated kick is up to $40 \text{ G}\cdot\text{m}$, which would deflect a 25 MeV beam approximately 50 mrad. Figure courtesy of A. Krasnykh, SLAC.

A traveling wave kicker with an effective field $2H$ produces an integrated kick strength

$$s_0[\text{T m}] \approx 2Hl = \frac{2E_{\text{max}}l}{ca}, \quad (3.1)$$

where E_{max} is the maximum pulse voltage at each plate, a is the half-aperture, and l is the kicker length. The beam deflection angle is then related to the beam's magnetic rigidity $B\rho$ [T m] = $3.33 \times 10^{-3} E$ [MeV],

$$\theta_k = s_0/(B\rho). \quad (3.2)$$

Using the numbers from Fig 3.2, the total kick strength is sufficient to move on orbit an electron bunch which has drifted ≈ 5 mm transversely from a septum magnet exit

10 cm upstream.

The peak powers for this kicker are high, although the average power is small: the peak power to the loads scale like $2E_{\max}/R \simeq 5 \text{ MW}$; the average power is multiplied by the duty cycle $\sim 10 \text{ ns} \cdot 60 \text{ Hz}$, for a total power of only $\simeq 3 \text{ W}$. The technical difficulty will likely be the pulsed power supply, although FET switches or other solid-state switching techniques have been developed that can handle the high peak power delivery.² Although the particular design for this ring is new, very high repetition rate designs with 7 ns rise-times and kV-level voltages are already in use [36].

Septum

The septum can be a Lambertson-style design which has the advantage of large coil windings since the field is perpendicular to the septum (Fig. 3.3). For compactness, the magnet can be positioned above the beam so that the injected bunch will drift down and be vertically kicked on to the closed-orbit axis. From the rigidity of the beam $B\rho$, a 7 kG field will bend a 25 MeV beam in a curvature $\rho \approx 12 \text{ cm}$. For a total angular deflection α , the physical magnet length l_m scales like $l_m = 2\rho \sin \frac{\alpha}{2}$. Using the above example, a 10 cm long magnet would bend the beam through an angle of 49° . The exact angle needed will likely be smaller but will depend on clearance issues and mechanical details of the transport line to ring layout.

The geometry of the injected beam can be seen from the end-view of the magnet in Fig. 3.3. The bunch begins to the right of septum on the far side of the gap and enters the magnet with a slight vertical downward trajectory (which is removed later by the kicker). The septum field moves the beam horizontally until it exits the magnet just above the ring closed-orbit trajectory. The remaining vertical offset is calculated from the beam clearances: 10σ of the circulating beam size, and a usual smaller 6σ on the injected beam. From the $5 \mu\text{m}$ emittance requirement and a β -function $\approx 1 \text{ m}$, $10\sigma \approx 3 \text{ mm}$, for an overall expected displacement of $\lesssim 5 \text{ mm}$ including some small septum thickness. This value is consistent with the kicker strength estimate for a

²Proof-of-principle experiments were conducted by A. Krashnykh (SLAC) for a drift step recovery diode technique as well as for an electromagnetic shock wave excitation on a ferrite-loaded coaxial cable, where both produced $\sim 10 \text{ ns}$ pulses at a few kV into a 50Ω load.

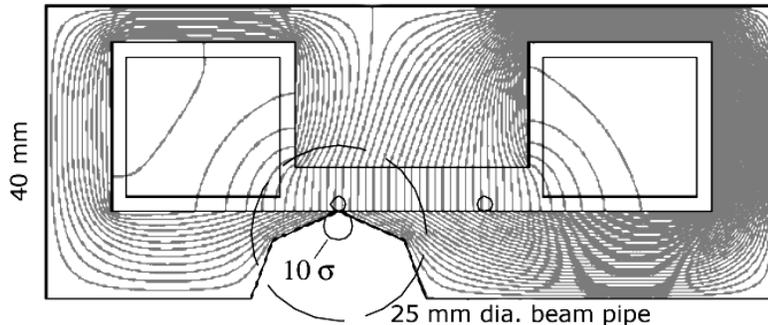


Figure 3.3: Lambertson septum. The gap distance is 6 mm and the field strength is ≈ 7 kG. Other details described in text. Figure courtesy of A. Krasnykh, SLAC.

septum to kicker separation of ~ 10 cm.

3.1.4 Beam Dynamics

Since the stored beam is thrown away each injection cycle (~ 10 ms or 10^6 turns), the only dynamics of interest must happen on relatively short timescales. In particular, radiation damping usually dominates electron storage ring behavior, but in this case the damping rates are \sim second (Table 3.1), including the laser radiative damping (see Sect. 2.3 pg. 20). Rather, the relatively low energy of the beam makes Coulomb, or intra-beam scattering (IBS) an important effect on emittance evolution. IBS was extensively studied at SLAC for low energy rings [87], and summarized below. Other single-bunch collective effects may lead to beam loss or beam emittance degradation; their relative thresholds can be estimated but not modeled until actual ring impedances are calculated. Therefore, a survey of the various instabilities needs to be checked before accepting any specific design of the ring. Resistive wall instabilities, microwave instabilities, and other collective effects can be well studied with computer models; experts in this field can very accurately predict a ring's performance before any hardware is built. Overall, the main advantage of this ring is that the long 1 cm bunch length reduces space-charge effects and provides modest peak currents that limit the strength of induced wakefields. Highlights are given below for a few of the more unusual aspects of the ring.

Chromaticity

A natural negative chromaticity causes tune spread and possible head-tail instability. For the lattice in Fig. 3.1, the maximum tune spread $\Delta\nu = \xi(\sigma_E/E)$ is < 0.03 , which is small. The head-tail instability is a function of transverse impedance Z_{\perp} , which should be modeled for the entire ring; Z_{\perp} can be estimated for resistive wall impedance, for example, which scales as b^{-3} where b is the pipe radius. Whether or not the growth rate is fast enough to need correction would need to be evaluated for the particular lattice and ring mechanical details.

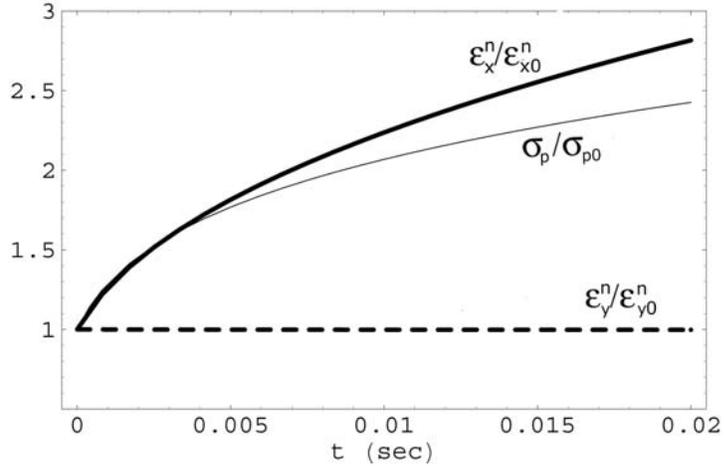
Intra-Beam Scattering (IBS)

Because of the relatively low emittance required in the circulating beam and a rather modest electron energy, intra-beam scattering [65] was identified as a dominant mechanism for emittance growth during millisecond timescales. IBS leads to diffusion in all three degrees of freedom through an exchange of energies between the longitudinal and transverse oscillations. The growth rate scales with intensity and is highly sensitive to the 6-D phase space of the beam.

Using the Bjorken-Mtingwa IBS formalism [8], local growth rates for the two transverse emittances and relative momentum spread can be calculated analytically as a coupled set of equations once the specified lattice functions and initial beam parameters are given. An example of such a calculation is shown in Fig. 3.4 for the proposed lattice (Fig. 3.1). The calculation does not include vertical dispersion or transverse coupling so the vertical emittance remains constant. From a practical standpoint, some vertical coupling would be advantageous to share the growth rate between both transverse degrees of freedom.

There was also an investigation about how to control IBS rates by deliberate lattice design. The local growth rates were calculated along the lattice and dominant terms identified. In the horizontal, for example, the growth rate is strongly correlated with $\mathcal{H}_x(s)$, where $\mathcal{H}_x = \beta_x \eta_x'^2 + 2\alpha_x \eta_x \eta_x' + \gamma_x \eta_x^2$ is expressed in terms of the usual lattice and dispersion functions. The highest growth rate occurs during the bend sections where the dispersion and small β -function waists locally increase \mathcal{H}_x . A

Figure 3.4: Normalized emittance and momentum spread growth due to IBS with initial parameters $\varepsilon_{x0}^n = \varepsilon_{y0}^n = 5\mu\text{m}$, $\sigma_{p0} = 10^{-3}$, $N_e = 6 \times 10^9$, and $\sigma_{s0} = 1\text{ cm}$. The RMS bunch length evolves as $\sigma_s(t) = \sigma_{s0}\sigma_p(t)/\sigma_{p0}$. Figure reproduced from Venturini [87].



better optimized lattice design may slightly improve the IBS growth rates further.

Even with the proposed lattice, the predicted IBS growth rates will cause an emittance degradation of a factor of 2–3 before the beam is refreshed. A factor of 2 may be a quite reasonable target when choosing a nominal re-injection rate. Since IBS scales with particle number, there is an optimization between beam current and emittance growth when considering the eventual output X-ray brightness.

Coherent Synchrotron Radiation (CSR)

Since the dipole magnets bend the beam through large angles, there is a possible bunch self-interaction via coherent synchrotron radiation [10]. The “overtaking length” is defined as

$$L = 2(3\sigma_s\rho^2)^{1/3}, \quad (3.3)$$

which is the length at which the bunch experiences steady-state interaction. For $\sigma_s = 1\text{ cm}$ and $\rho = 16.6\text{ cm}$, $L = 19\text{ cm}$ or angle $\theta = 65^\circ$. CSR has the practical consequence of causing an energy spread through the bends which might destroy the achromatic optics in the ring arcs. Shielding CSR by the beam pipe can be estimated by a parallel plate model where the vertical gap between plates h produces a shielding factor,

$$\eta = \sqrt{\frac{2}{3}} \left(\frac{\pi\rho}{h}\right)^{3/2} \left(\frac{\sigma_s}{\rho}\right). \quad (3.4)$$

For the nominal ring parameters and a gap of 2 cm, $\eta \approx 6$. The ratio of shielded to unshielded (free-space) CSR power is given by

$$P^{\text{sh}}/P^{\text{fs}} \simeq 4.2\eta^{5/6}e^{-2\eta}. \quad (3.5)$$

For $\eta = 6$, the ratio $P^{\text{sh}}/P^{\text{fs}} \approx 10^{-4}$, which is very well shielded. Transient effects can be higher than the steady-state estimates where even fields in the straight sections, before entering the bends, can influence the bunch [54]. For this reason, simulations of the ring are necessary to evaluate the required shielding, although the quick estimates above suggest that CSR should not pose a major problem.

3.2 Injection System

The injector is the most straight-forward system from a technical feasibility standpoint, although it contains the most hardware and cost compared to the ring and optical cavity systems. The system might consist of an rf photocathode electron gun, one or more linear accelerator section(s), and transport beam optics to match into the ring. The injector's function is to deliver \sim nC electron bunches at $\lesssim 5 \mu\text{m}$ transverse (normalized) emittance, with an energy spread $\lesssim 0.3\%$, at repetition rates between 30–120 Hz. Note the total beam power delivered to the ring (which is also the dumped beam power in the ring) is only \sim Watt.

RF Photocathode Gun

RF photocathode injectors have been under intensive development since the mid-80's and are now the preferred electron source for high-brightness applications. Typical design specifications are $\varepsilon_t^n \sim 1$ (π mm mrad) at 1 nC and they routinely produce experimental results of $\varepsilon_t^n = 2\text{--}3$ (π mm mrad). The usual metal cathodes are robust but their limited quantum efficiency often requires significant laser pulse energy. Semiconductor, or more exotic cathode materials offer much higher quantum efficiency but often do not perform well in the high-field rf environment.

The specifications for the ring emittance are well away from state-of-the-art performance of these rf photo-injectors. A conservative approach would be to use a standard Cu cathode which has a measured quantum efficiency $\approx 3 \times 10^{-4}$, requiring laser pulse energies of $\approx 150 \mu\text{J}$ at $\lambda = 266 \text{ nm}$ to produce a 1 nC bunch. Solid-state laser systems are commercially available at these power levels. There is also some freedom to choose an rf frequency, likely either L-band or S-band, depending on availability and cost of klystron sources, especially since the bunch length out of these injectors will be shorter than necessary to match into the ring. The optimization of the photo-injector design is then more one of cost and complexity rather than performance.

Linear Accelerator

One or more accelerator sections bring the few MeV beam up to the operating energy of 25 MeV. The rf frequency can be chosen—in conjunction with the rf photocathode—to leverage available klystron power sources, which are likely the dominant cost. Both L-band and S-band structures can provide robust accelerating gradients $\sim 20 \text{ MV/m}$ which is consistent with the overall scale of the ring. Accelerator structure design and manufacturing is a very mature technology that can be optimized for performance vs. cost.

Transport Line

The input beam must be optically matched to that of the ring at injection. The conceptual picture (Fig. 1.2 pg. 8) shows the injector system with a single dipole bend; however, in practice, the beam transport would likely be slightly more complicated to control dispersion. Transverse optics matching is straightforward; the only inherent mismatch between the injection beam and the beam in the ring is longitudinal. If possible, a bunch stretcher in the transport line could trade-off bunch length vs. energy spread, or else the shorter injected beam will quickly filament inside the ring. Whether these dynamics are detrimental to stability can be predicted with simulations, but the overall transport system requirements should be otherwise unremarkable.

Chapter 4

The Optical Storage System

The optical storage system would consist of an external enhancement cavity resonantly driven by a CW mode-locked laser. There are three major performance demands on the optical enhancement cavity: low internal losses, a narrow transverse waist at the interaction point (IP), and optical/mechanical stability. The drive laser has two major requirements: efficient coupling to the cavity eigenmode by transverse mode-matching and alignment, and tracking the central frequency of the cavity by maintaining an overall frequency stability within the cavity bandwidth. One mutual requirement is to match the laser pulse repetition rate to the free spectral range of the cavity, where both also match the electron storage ring circulation frequency.

An examination of these criteria is explored in this chapter, with emphasis on the technical requirements needed to fulfill the nominal laser-electron storage ring design.

4.1 The Power Enhancement Cavity

The issues of accumulating a large circulating power in an external cavity are well known for single-mode CW lasers. The added complication with a mode-locked laser is that the mode structure of the laser must match the mode structure of the cavity, or in time domain, the round-trip time in the cavity must very closely match the

repetition rate of the laser.¹ The following section describes a basic, two-mirror Fabry-Perot cavity in order to illustrate the various longitudinal and transverse issues, some of which are directly relevant to the experiments performed in Ch. 5.

4.1.1 The Fabry-Perot Cavity

The simplest optical cavity configuration is a two-mirror Fabry-Perot interferometer in which power is coupled through the backside of one partially transmissive mirror (Fig. 4.1). The mirrors have a spherical radius of curvature to provide stability for the fundamental transverse eigenmode of the cavity (TEM_{00}). The mirror radii, together with the cavity length, fully determine the waist size and position, the mirror spot sizes, and the inherent stability, *i.e.* sensitivity to alignment and transverse mode-matching errors. The mirrors' reflectivity, combined with transmission and other losses, determine the external coupling and cavity finesse—or expected gain. Higher finesse cavities mean longer fill times (lower cavity bandwidths) and correspondingly more stringent frequency stability of the driving source (discussed in Sect. 4.3.1).

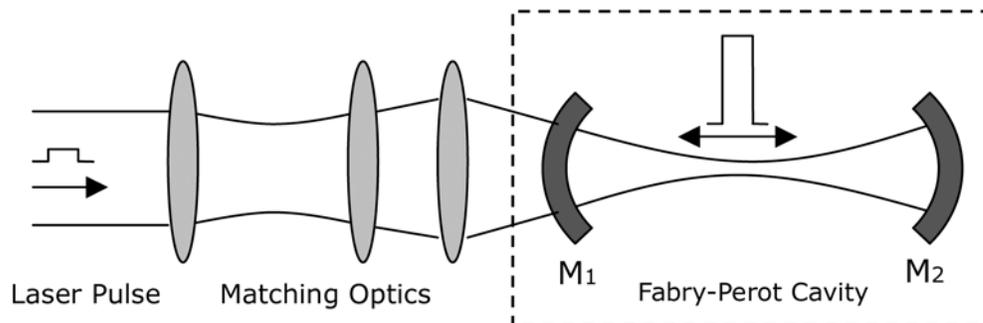


Figure 4.1: A basic pulse-stacking cavity geometry is a symmetric, standing-wave Fabry-Perot interferometer that has a round-trip circulation time equal to the inverse of the laser pulse repetition frequency and a central waist determined by the cavity eigenmode. Under steady-state, matched conditions, the power of the drive laser matches the losses of the cavity and results in an amplified circulating laser pulse with a gain $\sim \text{losses}^{-1}$.

¹or more generally they should be related by a harmonic, either for storing multiple pulses in a longer cavity, or having one pulse undergo multiple bounces in a folded cavity.

4.1.2 Cavity Performance at Resonance

General resonance properties of passive cavities are derived in Siegman [71, ch. 11]. The relevant cavity properties discussed below use “low-loss” cavity approximations and are valid for all experimental configurations in Ch. 5 as well. Dispersive effects, and other mode mismatches between the laser and cavity, are later described in Sect. 4.2.3.

Neglecting the requirements of the drive source for the moment, the most important cavity parameter necessary to achieve high circulating power is mirror quality—high reflectivity and high power handling. Mirrors are characterized by their reflectance R , transmissivity T , and losses L where $R+T+L = 1$ by energy conservation. The losses include both scattering and absorptive losses, the latter being less desirable as it can lead to thermal distortions of the mirror surface [47]. General power handling and manufacturing quality of multilayer mirrors is discussed separately in Sect. 4.1.4.

In terms of characterizing the performance of an optical cavity, a useful figure of merit is the “bounce number” b which is defined from the round-trip power loss in a cavity, $\propto e^{-1/b}$. The bounce number is just the e-folding decay time, similar to cavity Q , but measured in cavity round-trips. In the limit of small losses,

$$b = \frac{1}{T_1 + L_1 + T_2 + L_2}, \quad (4.1)$$

where subscript 1 always refers to the input (coupling) mirror. Other propagation losses—like scattering in air—are typically much smaller (\sim ppm) than mirror losses. These path losses can simply be added to $L_1 + L_2$ since, in practice, only the aggregate of non-transmissive losses can be reliably measured (see Section 5.2.3). The cavity finesse \mathcal{F} , usually defined in terms of reflectivity, is proportional to b and is useful in determining the cavity bandwidth $\Delta\nu_{\text{cav}}$ given the axial-mode interval between resonances, or “free spectral range” FSR,

$$\mathcal{F} \equiv \frac{\pi \sqrt[4]{R_1 R_2}}{1 - \sqrt{R_1 R_2}} \simeq 2\pi b \simeq \frac{\text{FSR}}{\Delta\nu_{\text{cav}}}, \quad (4.2)$$

where $\text{FSR} = c/2L$, L being the mirror-to-mirror separation distance. The gain, defined as the steady-state cavity pulse energy U_0 in units of incident pulse energy U_{inc} is

$$\text{gain} \equiv \frac{U_0}{U_{\text{inc}}} \simeq 4 \mathbb{T}_1 b^2. \quad (4.3)$$

The ‘‘impedance’’ match measures how well the incident field cancels the leakage field from the cavity in steady-state and is calculated from the net reflected field

$$\frac{\vec{E}_{\text{refl}}}{\vec{E}_{\text{inc}}} \simeq 1 - 2 \mathbb{T}_1 b, \quad (4.4)$$

where the square of this term determines the amount of power unable to couple into the cavity. This coupling then places an upper bound on the expected efficiency of the system. In a matched configuration, the input coupling equals the sum of all other losses in the cavity, $\mathbb{T}_1 = \mathbb{L}_1 + \mathbb{L}_2 + \mathbb{T}_2 = 1/2b$, yielding no net power reflection. From a practical standpoint, the coupling mismatch has a broad minimum since the reflected power varies quadratically with \mathbb{T}_1 near match.

As an example, a designed cavity enhancement of 10,000 requires a total cavity loss, including mirror transmission losses, of 200 ppm of which 100 ppm is ideally the transmission \mathbb{T}_1 . The resulting cavity finesse is $\approx 30,000$, well below what is achievable with low-loss mirrors [20].

4.1.3 Cavity Alignment Sensitivity

The cavity geometry is commonly characterized by the g parameters,

$$g_i = 1 - \frac{L}{R_i} \quad i = 1, 2 \quad (4.5)$$

where L is the cavity length and R_i are the mirrors’ radius of curvatures. The stability condition $0 \leq g_1 g_2 \leq 1$ is analogous to the stability condition required in electron beam optics for a periodic transformation matrix.²

The g parameters are useful in estimating the expected alignment sensitivity of

²the analogy in accelerator physics for this cavity is a simple ‘FOFO’ lattice where the thin lens focus $f = R/2$.

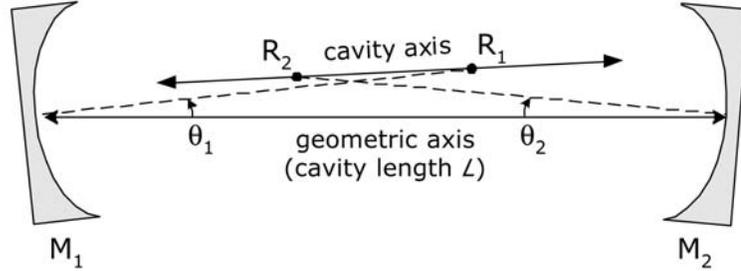


Figure 4.2: Optical resonator misalignment geometry for estimating angle and offset uncertainty of the optical axis with respect to the geometric axis.

the cavity optical axis, defined as the line connecting the centers of curvature for each mirror, with respect to a defined geometric axis, *i.e.* the axis running through the center of each optic (Fig. 4.2). The analysis of tilt and offset errors of the cavity mirrors has been studied in detail [37, 38]. The uncertainty in the cavity axis slope, $d\alpha$, and offset, da , referenced from the input mirror can be expressed as

$$d\alpha^2 = \frac{d\theta_1^2 + d\theta_2^2}{(1+g)^2}, \quad \text{and} \quad da^2 = L^2 \left[\frac{g^2 d\theta_1^2 + d\theta_2^2}{(1-g^2)^2} \right] \quad (4.6)$$

where L is the cavity length, and $d\theta_1$ and $d\theta_2$ are the uncertainties in the angles of mirrors M_1 and M_2 . This relation applies to each transverse axis. The actual alignment requirements, and the consequences of misalignment, are discussed in Sect. 4.2.1.

As an example of alignment sensitivity, the experimental cavities tested in Ch. 5 have a cavity $L = 1.9$ m and $R = 1.0$ m which gives $g = -0.90$. The relative uncertainty in the optic axis is then $d\alpha^2 = 100(d\theta_1^2 + d\theta_2^2)$ and $da^2 = 43d\theta_1^2 + 53d\theta_2^2$. This implies that the precision in aligning the mirror tilt translates to roughly an order-of-magnitude worse alignment precision in setting the cavity optic axis.

A two-mirror resonator designed to integrate in the electron ring would be extremely sensitive to misalignment. Such near-concentric cavities ($g \approx -1$) have been studied for FELs [18]. For a symmetric cavity, the waist size is given by

$$w_0^2 = \frac{\lambda L}{2\pi} \sqrt{\frac{1+g}{1-g}}, \quad (4.7)$$

which implies that the distance to instability, δ , where $g = -1 + \delta$, scales like w_0^4 . For a nominal design, a $60\ \mu\text{m}$ waist with $\lambda = 1\ \mu\text{m}$ and $L = 1.7\ \text{m}$, $g = -0.9997$. Attempting to focus to a $30\ \mu\text{m}$ waist instead of a $60\ \mu\text{m}$ waist would imply another order-of-magnitude more sensitivity to misalignment errors. The alignment precision required for a two-mirror cavity would require some active servo control on the mirrors [2], as well as careful manufacturing and measurement of the mirror curvature [84,92]. Another possibility is to design a more complicated cavity which affords narrow waists but is naturally more stable [57]. The parameters chosen for the X-ray source design are demanding but should fall within the technical solutions available.

4.1.4 Mirrors

Progress in coating technologies over the last decade has enabled the routine manufacture of extremely low-loss (\sim ppm) optics. Fabry-Perot cavities can be easily made with over 150,000 finesse—corresponding to a total cavity loss of 40 ppm, which includes the input transmission [20]. Scattering and absorption losses are typically under 10 ppm per mirror for commercially available optics [81, 82], while research efforts for LIGO has recently produced optics with losses under 1 ppm.³

These optics are multi-layer dielectric coatings produced by ion beam deposition on top of super-polished substrates. The thin films are quarter-wave stacks of alternating high and low indexes of refraction and are made with materials chosen for low mechanical stress, physical robustness, and low absorption. The various metal oxides typically used in these coatings are well characterized and known within the industry [5].

The optical cavity requires two additional attributes, however, besides low-losses: high power-handling, and radiation resistance. Fortunately, the properties which make these optics low-loss also make them more resistant to damage. Optical damage limits are usually given for either pulsed power (J/cm^2) or CW power (MW/cm^2) depending on the source laser and corresponding damage mechanism. The repetition rate of CW mode-locked lasers is so high as to make them essentially CW sources with

³private communication with Gary DeBell, MLD Technologies.

respect to any materials effects. The nominal guideline for maximum power density is 1 MW/cm^2 , although this value has been experimentally exceeded (2.1 MW/cm^2) without any apparent loss of performance [84]. The narrow waist in the cavity will naturally result in large spots on the mirrors which would keep the optical power density within these limits (see next section).

Radiation damage of dielectric coated mirrors is less understood. The majority of published work has studied radiation damage to high reflectance mirrors used in FEL cavities, which are often very harsh environments. The electron beams required to drive FELs emit a broad spectrum of synchrotron radiation and deposit significant amounts of heat on the mirror surface. The UV to X-ray radiation has led to several accounts of contamination resulting from breakdown of hydrocarbons on the mirror surface [33–35, 94]. However, the carbon contamination has also been successfully removed by the use of RF-induced oxygen plasma [94]. Intense heating, on the other hand, has led to direct ablation of the coatings, which is irrecoverable [95]. These studies do share the fact that ion-beam coated deposition is by far more robust than other conventional coating techniques, likely due to low contamination and high packing density of the multi-layers. Other studies have measured the effect of large pulsed doses of soft X-rays on optical surfaces [12, 25]. Here, the damage mechanism is again intense heating due to radiation absorption in the high index of refraction material (containing heavier Z elements) just below the top layer of typically low absorbing silica.

For the nominal source design, the X-ray radiation incident on the cavity mirrors is several orders-of-magnitude lower in intensity, and of much narrower bandwidth, than the undulator radiation in storage ring FELs. The most likely concern will be whether X-rays contribute to a long-term degradation of mirror reflectivity through ionization effects in the multilayers (*e.g.* atomic lattice displacements). Since the source spectrum is quasi-monochromatic, even much of this X-ray absorption might be minimized with careful selection of dielectric materials that have low absorptions over these X-ray energies (see Sect. 2.5).

One final note is an encouraging anecdotal report from the Compton polarimeter at JLab [27]. A high-finesse cavity made with ion-beam coated dielectric mirrors

(≤ 16 ppm losses) showed no signs of degradation after 4 months of running a 3 GeV ($40 \mu\text{A}$) electron beam just 5 mm away from the mirrors.

Mirror Size

The mirror size must be chosen sufficiently large enough to contain the TEM_{00} mode with little diffraction losses while ensuring large enough spot sizes to keep the power density under $\sim 1 \text{ MW}/\text{cm}^2$.

The spot size $w(z)$ on the mirrors a distance z from the waist is directly related to the waist size, or depth of focus

$$z_R = \frac{\pi w_0^2}{\lambda}, \quad \text{and} \\ w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2} \quad (4.8)$$

where z_R is called the Rayleigh range (analogous to β^* in particle beam optics). From this scaling, a z_R of 1 cm ($w_0 = 60 \mu\text{m}$) will produce a spot size $w = 4.5 \text{ mm}$ on a mirror 0.75 m away. For a 1 mJ pulse circulating at 100 MHz repetition rate, the power density on these mirrors is approximately $160 \text{ kW}/\text{cm}^2$, still well below any optical damage threshold.

The required overall mirror size scales with the spot size with a factor determined by allowable diffraction losses. For a Gaussian beam, the diffraction losses are $\exp[-2a^2/w^2]$ where a is the mirror radius. To ensure < 1 ppm diffractive loss, $a \geq 3w$. In the example above, a diameter of 27 mm suffices; standard 2 in. optics or slightly smaller custom optics are straightforward solutions.

4.2 Driving the Enhancement Cavity

The resonant build-up of pulsed power in an enhancement cavity requires that the drive laser ideally matches both the transverse cavity eigenmode as well as the group and phase velocities of the round-trip pulsed fields circulating in the cavity. The following sections examines how closely the mode-locked laser pulses are expected to

meet these requirements.

4.2.1 Steering and Mode-Matching Errors

In general, an arbitrary input beam may be expressed as an infinite sum of cavity eigenmodes. Proper alignment of the input beam to the cavity is defined when an incident Gaussian beam couples power only to the fundamental eigenmode of the cavity and to no other higher-order modes. Several authors discuss both general analytic [6] or numerical treatments [37] of misalignment effects originating from either geometric alignment or mode-matching errors. A brief summary of cavity modes is reviewed first, followed by a description of matching errors and their relative sensitivities.

Transverse Modes

The cavity modes are conveniently described in cartesian coordinates by orthogonal Hermite-Gaussian functions [71]. An important consequence of these functions is that a mode-dependent (Guoy) phase shift affects the spacing of their resonant frequencies. For a general TEM_{mnq} mode, the axial modes (labelled by the longitudinal mode number q) are spaced at the cavity FSR ($= c/2L$), but higher-order transverse modes are spaced by a Transverse Mode Range (TMR)

$$\text{TMR} = \frac{1}{\pi} \cos^{-1} (\pm \sqrt{g_1 g_2}) \text{ FSR}. \quad (4.9)$$

A high-order mode TEM_{mnq} appears at a frequency ν_{mnq} ,

$$\nu_{mnq} = \nu_{00q} + (m + n)\text{TMR}. \quad (4.10)$$

This last relation is written to emphasize the fact that a cavity resonant at a fundamental mode frequency, ν_{00q} , will in general *not* be resonant for other higher order modes. Even for the designed near-concentric two-mirror cavity, the TMR ~ 500 kHz requires just a finesse of 200 to ensure a cavity bandwidth smaller than the next higher-order mode resonant frequency.

The result of this analysis is that for the source design, as well as for the experiments in Ch. 5, the effect of transverse matching errors is simply reduced coupling efficiency: higher-order modes will not be resonant with the fundamental mode and so are promptly reflected at the cavity input under normal operating conditions. This fact can be used to measure the effective cavity TEM₀₀ content of the input beam experimentally (see Sect. 5.2.3).

Matching Parameters

There are four parameters to match in each transverse plane: two from alignment—tilt, α , and offset, a , of the incident beam with respect to the cavity axis, and two from mode-matching—conveniently parameterized by waist size, w'_0 , and axial displacement, Δz_0 . For the case where only small mismatches are expected, Anderson [2] describes an intuitive approach to estimate mode couplings using a simple linear perturbation analysis. This analysis reveals that any small misalignments couple to lower-order modes with a phase dependent on the type of misalignment. A summary of these results is reproduced in Table 4.1.

A more general numerical study of misalignments by Haar [37] suggests that each of these matching errors fits very well to Gaussian functions. The perturbation analysis, on the other hand, is linear in fields so the resulting matching errors are fit to a parabolic function. For coupling mismatches down to 75% off peak, however, these expressions agree in principle to within 1%.

For the two-mirror source design cavity, the most sensitive matching will likely be tilt and offset misalignment given the difficulty in defining the cavity optical axis (Sect. 4.1.3). For instance, the cavity axis angle might be determined with good precision using the transmitted power through the cavity, but from Table 4.1, the input beam offset should be less than the cavity waist size, $w_0 = 60 \mu\text{m}$, which may be difficult to evaluate given the large spot sizes on the mirrors (where the measurements are done).

Table 4.1: Coupling coefficients for a mismatched beam driving a Fabry-Perot cavity (reproduced from Anderson [2]).

Misalignment	Term	Coupling Coef.	Mode	Phase (deg)	Frequency
cavity axis offset	a_μ	a_μ/w_0	$U_1(\mu)$	0	ν_0
cavity axis tilt	α_μ	$\alpha_\mu(\pi w_0/\lambda)$	$U_1(\mu)$	90	ν_0
mode waist size	w'_0	$w'_0/w_0 - 1$	$V_1(r)$	0	$2\nu_0$
mode waist position	Δz_0	$\lambda\Delta z_0/2\pi w_0^2$	$V_1(r)$	90	$2\nu_0$

The coupling coefficients describe the excitation strength of a higher order mode given in the next column. For example, $\Psi(x) = AU_0(x) \Rightarrow A[U_0(x) + (a_x/w_0)U_1(x)]$ where an initial beam of the lowest Hermite-Gaussian $U_0(x)$ excites some of the next higher-order mode $U_1(x)$ (TEM₁₀); $\mu = x$ or y , and ν_0 is one TMR from the fundamental resonance (Eq. 4.10). For mode mismatches, the analysis is performed in polar coordinates where the expansion is in Laguerre-Gaussian modes V_0 and V_1 . Notice there is enough phase and frequency information in the excited modes to determine the sources of each error term; this feature also makes a dynamic feedback servo possible.

4.2.2 Axial Mode Matching

A steady-state resonance condition occurs when the circulating pulse returns after one cavity round-trip to the coupling mirror such that its fields coherently add to an input drive pulse train just enough to compensate the lost round-trip energy. This condition for resonance, however, depends on a timely superposition of fields, not the envelope function, so that a collection of resonances can be described not only at the matched pulse envelope repetition frequency, but also at every $n\lambda$ displacement of this matched length. Physically, this corresponds to a series of resonances as the envelope peak “slips” integer wavelengths, or cavity axial modes. The following models describe the circulating cavity pulse under these general conditions of resonance.

Analytical Model

Taking the case of a matched cavity on resonance, the input mirror transmission T equals the combined other losses in one round-trip of the cavity. Let the incident pulse be described as a Gaussian field envelope of length σ (in \hat{z}) where a constant

envelope shift of ψ is allowed per round-trip (corresponding to an axial mode shift from peak). The cavity pulse envelope $\hat{E}(z)$ is built-up by a progression of n incident pulses. The few first terms are

$$\begin{aligned}
\hat{E}(z)_{n=1} &= \sqrt{\mathbb{T}} e^{-\frac{z^2}{2\sigma^2}} \\
\hat{E}(z)_{n=2} &= \sqrt{\mathbb{T}} \left[(1 - \mathbb{T}) e^{-\frac{(z-\psi)^2}{2\sigma^2}} + e^{-\frac{z^2}{2\sigma^2}} \right] \\
\hat{E}(z)_{n=3} &= \sqrt{\mathbb{T}} \left[(1 - \mathbb{T})^2 e^{-\frac{(z-2\psi)^2}{2\sigma^2}} + (1 - \mathbb{T}) e^{-\frac{(z-\psi)^2}{2\sigma^2}} + e^{-\frac{z^2}{2\sigma^2}} \right] \\
&\vdots \\
\hat{E}(z)_n &= \sqrt{\mathbb{T}} \sum_{k=1}^n (1 - \mathbb{T})^k e^{-\frac{(z-k\psi)^2}{2\sigma^2}}
\end{aligned} \tag{4.11}$$

where the input peak field is normalized to unity. Numerically, such sums are straightforward to calculate and can be generalized to include non-linear phase shifts or other dynamics [37].

A closed-form solution is possible for the particular series described in Eq. (4.11) as $n \rightarrow \infty$ by casting it in integral form

$$\hat{E}(z) = \sqrt{\mathbb{T}} \int_0^\infty e^{-\frac{(z-k\psi)^2}{2\sigma^2}} e^{-k\mathbb{T}} dk \quad \text{where } \mathbb{T} \ll 1. \tag{4.12}$$

Solving this integral,

$$\hat{E}(z) = \sqrt{\mathbb{T}} \sqrt{\frac{\pi}{2}} \frac{\sigma}{\psi} \exp\left(\frac{\mathbb{T}(\mathbb{T}\sigma^2 - 2\psi z)}{2\psi^2}\right) \left(1 + \operatorname{erf}\left(\frac{\psi z - \mathbb{T}\sigma^2}{\sqrt{2}\sigma\psi}\right)\right). \tag{4.13}$$

This function describes the steady-state pulse envelope in z (or time), examples of which are plotted in Fig. 4.3. To find the expected power coupling efficiency, c_p , of this steady-state solution, the integrated energy of this pulse can be compared to that of the matched case ($\psi = 0$)

$$c_p = \frac{\int_{-\infty}^{\infty} \hat{E}^2(z, \psi) dz}{\int_{-\infty}^{\infty} \hat{E}^2(z, 0) dz}. \tag{4.14}$$

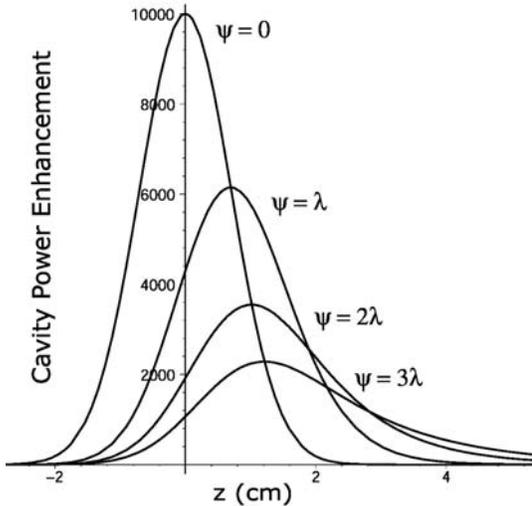


Figure 4.3: Analytical prediction of cavity power vs. axial mode slip. This example is for $\lambda = 1 \mu\text{m}$ and E-field envelope $\sigma = 1 \text{ cm}$ drive pulse with a nominal gain of 10,000 at peak. The pulse form is given by the square of Eq. (4.13) with a power coupling calculated from Eq. (4.14). The coupling efficiencies c_p are 100%, 75.8%, 54.6%, and 42.0%, respectively.

Note that this coupling efficiency gives a higher value than that predicted purely by measuring ratios of peak powers (slip vs. no slip) since the effective pulse shape widens as well. Predicting peak powers is still useful for diagnostics, however, since this is what photodetectors measure experimentally. Further study was performed with a numerical model to confirm these analytic predictions and illustrate the dynamic behavior of the cavity.

Numerical Model

A very simple but flexible simulation of the cavity fields can be described by iteratively calculating the forward and reflected field envelope on both sides of the input coupler. Fig. 4.4 depicts the fields and their relations applicable at resonance. These calculations apply to any general resonator although the two-mirror cavity (as tested in Ch. 5) is the main study. Stepping through this algorithm “bounce-by-bounce” reveals the dynamic fill-time behavior of the resonator.

Qualitatively, there are a few interesting features for small envelope timing slips (or advances). First, a slightly longer cavity creates a steady-state peak delayed after the input drive pulse, but the circulating pulse still remains very close to the Gaussian shape; the reflected fields from the cavity then show a characteristic shape (sketched as \mathbf{B}_n in Fig. 4.4) which looks like the derivative of a Gaussian pulse (since it’s two

Iteration of pulse waveforms:

$$\mathbf{B}_n = \mathbf{A}r - \mathbf{C}_n t$$

$$\mathbf{D}_n = \mathbf{C}_n r + \mathbf{A}t$$

$$\mathbf{C}_{n+1} = \mathbf{D}_n p; z \rightarrow z - \psi$$

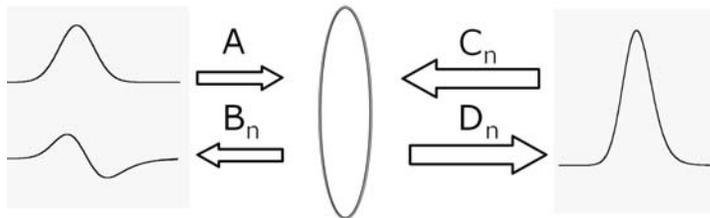


Figure 4.4: Field relations for dynamic evaluation of power build-up in a resonator at the coupling interface. Input \mathbf{A} represents a train of Gaussian shaped pulse fields. The values r and t are the reflection and transmission (field) coefficients of the coupling mirror, *i.e.* $t^2 = T_1$ and $r^2 = 1 - T_1$. For the two-mirror cavity case, the reflection coefficient of the output mirror is added to all other path losses so that $p^2 = 1 - T_2 - L_1 - L_2$. The algorithm is conveniently implemented in LabVIEW where the field relations are evaluated point-by-point on waveforms that are viewable dynamically as the simulation runs. To avoid waveform sampling effects, the calculation of arbitrary small displacements, ψ , is accomplished by adding a scaled numerical derivative of the circulating pulse to itself.

Gaussian functions slightly offset in time subtracted from each other). Another interesting feature is that the fill-time of the cavity changes character. On the peak mode, the fill-time behaves exponentially just as CW source, but off this mode, the peak gain abruptly saturates and reaches steady-state at a time much sooner than at peak (Fig. 4.5). By resonating at further axial mode slips, not only does the power coupling efficiency decrease but the effective cavity bandwidth changes as well (see Fig. 4.6). Looking ahead to Ch. 5, this feature was advantageous experimentally since it provides a means to vary the cavity bandwidth in small steps.

In terms of the intended source design, the results of these models suggest that operation is optimal at the peak axial mode, which is well defined but will undoubtedly require some slow feedback mechanism to track it. The peak circulating field, although easy to measure, is not as good a measure of performance as the power coupling efficiency since the value of interest is the circulating pulse energy. For the case of a 10,000 cavity gain, the results of Fig. 4.3 show a drop of 25% in coupling power for a slip of one axial mode off peak. Although only discrete envelope shifts associated with axial modes have so far been addressed, the tools developed in this section allow for arbitrary envelope phase slips, which is the topic of the next section.

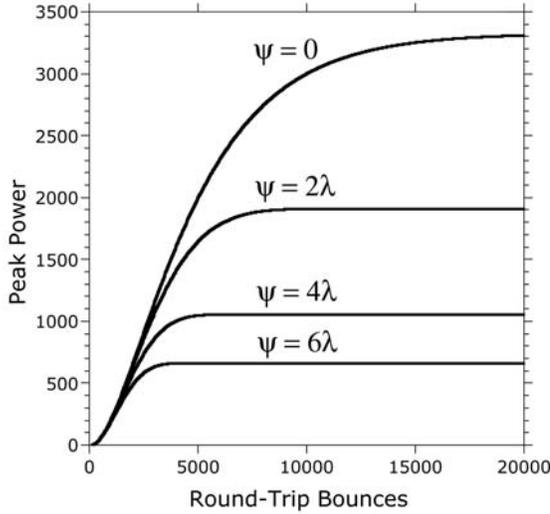


Figure 4.5: Dynamic peak power build-up in a resonator for a few axial mode slips. Although this example uses parameters relevant for experimental cavities described in Ch. 5, the dynamics illustrate the character of how power reaches steady-state in the resonator. These results are for a matched cavity with $T = 300$ ppm, $FSR = 80$ MHz, $\lambda = 1 \mu\text{m}$, and $\sigma = 0.6$ cm. Peak gain is ≈ 3300 and the nominal cavity bandwidth, $\Delta\nu_{\text{cav}} = 7.6$ kHz.

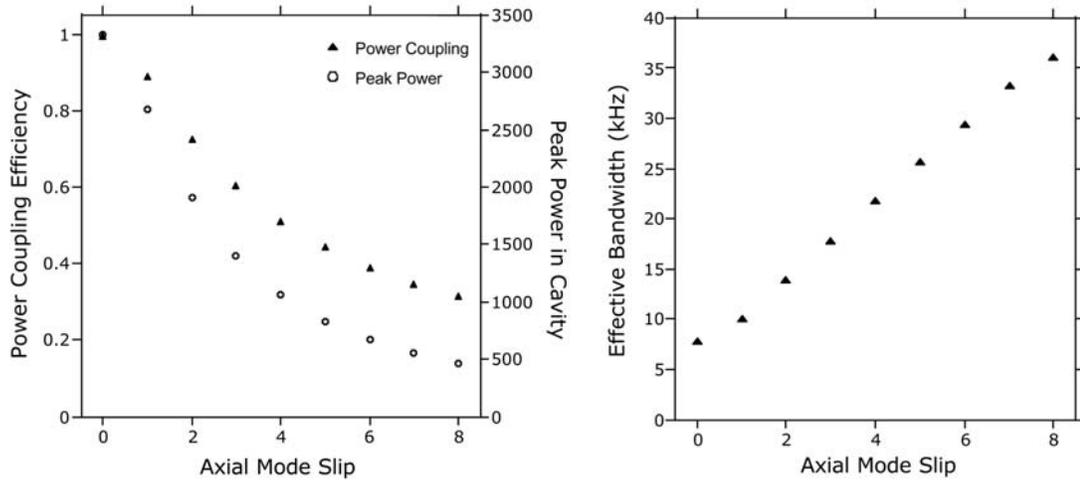


Figure 4.6: Power coupling efficiency and effective bandwidth vs. axial mode slip. The figure on the left contrasts the actual power coupling vs. peak power measurement for a series of integer λ envelope slips using the cavity parameters from Fig. 4.5. If either the pulse length or cavity finesse is known, the other can be solved. (Such a measurement of the laser pulse length was experimentally confirmed in Fig. 5.15, pg. 90.) The figure on the right measures the effective bandwidth of the cavity by taking curves like those in Fig. 4.5 and defining the fill-time ($\approx 6\tau$) as the time to achieve 90% of steady-state power. This tunability of the cavity bandwidth was an important experimental tool in measuring tracking stability of the laser center frequency to that of the cavity (see Sect. 5.3.3).

4.2.3 Dispersion Matching

The optical materials in mode-locked lasers generally exhibit a difference in phase velocity v_ϕ vs. group velocity v_g . For every round-trip in the laser cavity there is a cumulative phase offset between the envelope and carrier of the optical pulse given by

$$\Delta\phi_{RT} = \frac{2\pi}{\lambda} \int_0^L [n_g(z) - n(z)] dz = \frac{\omega^2}{c} \int_0^L \frac{dn(z)}{d\omega} \quad (4.15)$$

where L is the optical length of the laser cavity, $n = c/v_\phi$ and $n_g = c/v_g$ are the refractive and group indexes, and ω the angular frequency [39]. The total phase slip $\Delta\phi_{RT}$ can be quite large; for example, $\Delta\phi_{RT} \sim$ several hundred cycles for a few mm Nd:YVO₄ crystal (estimated using the Sellmeier coefficients at $\lambda = 1 \mu\text{m}$). However, the detectable carrier-envelope slip is only the fractional part,

$$\Delta\phi_S = \Delta\phi_{RT} \bmod 2\pi. \quad (4.16)$$

A frequency f_S can be defined for the rate at which the phase offset changes per round-trip time in terms of the laser repetition rate f_{rep}

$$f_S = \frac{\Delta\phi_S}{2\pi} f_{\text{rep}}. \quad (4.17)$$

For a static phase offset, this additional frequency term just shifts the comb of modes, uniformly spaced at f_{rep} , by some fixed frequency f_S . The fact that there are only two degrees of freedom in the laser frequency comb has enabled the development of novel methods in precision optical metrology [21, 80]. Dynamic phase offsets, for instance from thermal variations or acoustic-mechanical coupling, add to the overall frequency noise of the laser (see Sect. 4.3.1).

The result of dispersion in picosecond lasers is usually not a concern, although for very short pulses (< 10 fs), envelope phase stabilization is important for studies that are sensitive to the “absolute” carrier envelope phase [28]. Picosecond laser pulses, however, can be sensitive to these dispersion effects when used to drive a high-finesse cavity. Ideally, not only the cavity lengths are matched, but the dispersion of both

the laser and the external cavity are also matched such that their frequency combs are identical at steady-state. Using the estimate discussed above for the magnitude of the laser dispersion, the fractional phase $\Delta\phi_S$ is essentially fixed but random⁴ ($-\pi < \Delta\phi_S \leq \pi$). The dispersion in the external cavity, on the other hand, is easier to estimate; the only optical materials in the passive cavity are the multi-layer mirrors where the quarter-wave (QW) stack produces a highly linear relation between the group and phase velocities over the central bandwidth of the mirror [78]. The reflection phase from a QW stack ($\lambda_0/4$) near the central frequency $\omega_0 = 2\pi c\lambda_0$ equals

$$\varphi = \pi \left(1 - \frac{n_0}{n_H - n_L} \frac{\omega - \omega_0}{\omega} \right), \quad (4.18)$$

where the refractive indexes of the high and low layers are n_H and n_L (from a media n_0), and $\omega - \omega_0 \ll \omega$. The last assumption is valid for picosecond laser pulses where the typical bandwidth of the gain media is ~ 1 nm at $1 \mu\text{m}$, and the QW stack bandwidth scales as

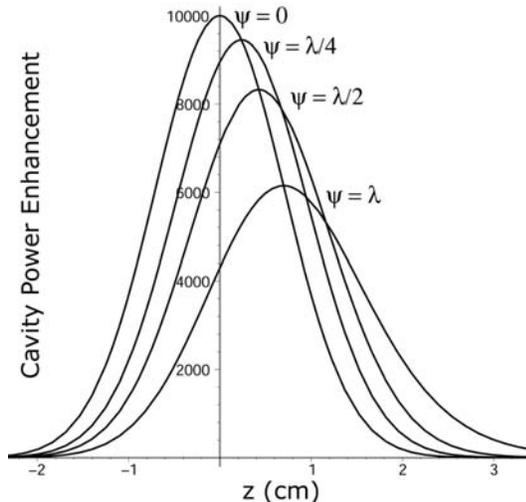
$$\left(\frac{\Delta\omega}{\omega_0} \right)_{\text{QW}} = \frac{2}{\pi} \sin^{-1} \left(\frac{n_H - n_L}{n_H + n_L} \right) \quad (4.19)$$

which is typically around 10% for commonly used dielectric coatings [58]. The slope $d\varphi/d\omega$ is constant and negative. The sign just expresses the fact that there is an additional phase lag beyond that predicted from the physical separation of mirrors since the pulse energy must penetrate somewhat into the QW layer stack. Besides this apparent lengthening of the cavity, the dispersion in the external cavity does not contribute any more to the physics of the circulating pulse. Importantly, there is no non-linear phase dependence on frequency which would cause a temporal widening of the pulse envelope. The dispersion ‘‘match’’ is now just a simple phase carrier-envelope offset.

How this predicted carrier-envelope offset affects the gain performance can be estimated using the same analysis for axial mode shifts in the last section. Instead of discrete envelope slips of $\psi = n2\pi$, the dispersion mismatch would appear as a shift of $|\psi| \leq \pi$. Taking the worst case scenario of $\psi = \lambda/2$ together with the nominal design

⁴‘fixed’ meaning the pulse-to-pulse variation is likely negligible although thermal or environmental changes will systematically affect the modulo- 2π ‘randomly’ selected leftover phase slip.

Figure 4.7: Analytical prediction of cavity power vs. dispersive carrier-envelope phase slip. This example uses the same parameters as in Fig. 4.3 for $\lambda = 1 \mu\text{m}$ and field $\sigma = 1 \text{ cm}$ drive pulse with a nominal gain of 10,000 at peak. The pulse form is given by the square of Eq. (4.13) with a power coupling calculated from Eq. (4.14). The coupling efficiencies c_p are 100%, 97.1%, 90.5%, and 75.8%, respectively. For the case $\psi = \lambda/2$, the effective pulse width also widens by roughly 9%.



parameters, the effective power gain only drops by 10% as illustrated in Fig. 4.7. Since the power coupling efficiency is less sensitive to carrier-envelope offsets near zero, the effect of this dispersion mismatch is likely to be quite small for these parameters in practice. Higher gain cavities—or shorter optical pulses—may require some kind of tunable dispersion control in the laser. Such control is an added complexity but not technically difficult, and is often implemented in femtosecond lasers by using standard prism sequences [70] or more sophisticated modulation techniques [88].

4.3 Laser Frequency Stabilization

For the drive laser to remain resonant with the external cavity, the laser central carrier frequency must track the cavity resonant frequency (“achieve lock”) by reducing the free-running frequency noise of the laser to some fraction of the external cavity bandwidth. A powerful technique to do this is Pound-Drever-Hall (PDH) stabilization [24] and is well described in a recent paper by Black [9]. A short overview of the technique is given below after a brief summary of CW mode-locked laser performance issues and a review of feedback control modeling.

4.3.1 Characteristics of CW Mode-Locked Lasers

Diode-pumped, solid-state lasers provide very stable and reliable sources for high-power, near-IR, picosecond optical pulses. CW, passively mode-locked Nd:YAG or Nd:YVO₄ (vanadate) lasers are commercially available that produce more than 10 watts of average power at 1064 nm with pulse widths as small as 7 ps in the TEM₀₀ mode. To avoid some dispersion effects due to the short pulse (see 4.2.3) an internal filter may be added to reduce the laser cavity bandwidth in order to produce long pulse widths without sacrificing beam quality. The mode-locked repetition frequency typically operates between 70–200 MHz and matches the expected design range for the electron storage ring circulation frequency.

In frequency domain, the number of CW modes contributing to the pulse train is given by the gain bandwidth—which is inversely proportional to the pulse width—divided by the repetition rate of the laser. For laser pulse widths in the 10’s of picoseconds operating at repetition rates near 100 MHz, there are typically several hundred modes within the bandwidth of the laser. The modes are uniformly spaced across this band (down to a part $\sim 10^{17}$ [80]), but may have a comb offset given by the dispersion within the laser (see Sect. 4.2.3).

The lowest achievable linewidth, or spectral bandwidth, of each mode is usually described by the Schawlow-Townes limit and is a measure of the spectral broadening due to quantum noise fluctuations (spontaneous emission). This limit is noticeable in only extremely stable lasers that have sub-Hz linewidths. In practice, the free-running frequency noise of solid-state lasers is dominated rather by “technical noise”—acoustic-mechanical environmental noise coupling to the laser cavity optics or gain media. The optical path length changes induced by this noise is manifested as frequency jitter on the entire mode spectrum.

In order to implement a frequency stabilization feedback, there must be a mechanism to adjust the laser cavity optical path in response to the measured noise. One common technique is adding a piezo driven mirror or other high-bandwidth actuator like an electro-optic crystal in the laser cavity itself.

4.3.2 Frequency Stabilization Feedback Model

Frequency stabilization means using an active feedback control system to suppress the laser's free-running frequency noise: the feedback system is composed of a discriminator signal (*e.g.* the free-running frequency of the laser compared to a reference frequency), which is then conditioned through a compensator circuit and applied back to the laser plant through an actuator. The basic block diagram, including noise sources, is illustrated in Fig. 4.8 [1].

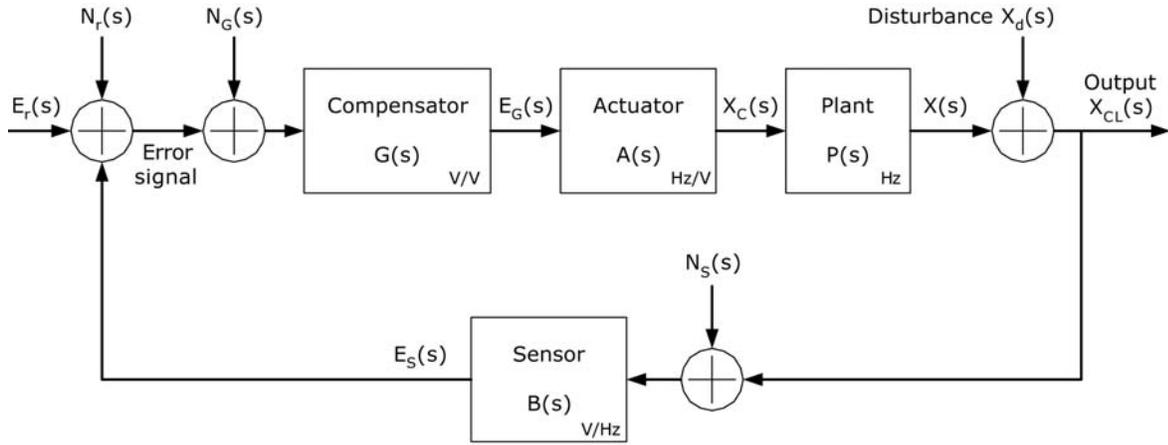


Figure 4.8: Frequency stabilization feedback system block diagram. Components are represented by their transfer functions and signals by their Laplace transforms. From the laser plant $P(s)$, the free-running frequency is $X(s)+X_d(s)$ where $X_d(s)$ represents the noise terms, while $X_{CL}(s)$ is the controlled output under closed loop conditions ($X_C(s) \neq 0$). The sensor $B(s)$ converts a frequency measurement to volts $E_S(s)$, with noise $N_S(s)$, which then is summed with a reference E_S , and noise $N_r(s)$, to produce an error signal. A compensator or servo $G(s)$ amplifies and filters the error signal and applies it to an actuator $A(s)$ which converts the voltage signal $E_G(s)$ back to frequency. The result is a correction signal $X_C(s)$ in Hz back to the laser plant $P(s)$.

The performance of the feedback loop can be estimated using some simplifying assumptions. The open-loop transfer function $L = BGAP$ defines the basic character of the system. Solving for the closed-loop output using the block diagram in Fig. 4.8,

$$X_{CL} = \frac{(E_r + N_r + N_G)}{B} \frac{L}{(1 + L)} - N_S \frac{L}{(1 + L)} + (X + X_d) \frac{1}{(1 + L)}. \quad (4.20)$$

Disregarding extra noise terms for the moment, the ratio of the sensor output BX_{CL} and control input E_r defines the closed-loop transfer function

$$\frac{BX_{\text{CL}}}{E_r} = \frac{L}{1+L} \simeq 1 \text{ for } |L| \gg 1. \quad (4.21)$$

For high open-loop gain $|L| \gg 1$, Eq. (4.20) reduces to

$$X_{\text{CL}} \simeq \frac{1}{B}(E_r + N_r + N_G) - N_S + \frac{(X + X_d)}{L}. \quad (4.22)$$

These relations show that for large open-loop gains, the error signal suppresses the free-running noise by the magnitude of the loop gain L , and that the noise terms from the sensor, reference, and compensator stages add with a weighting given by the sensor gain B . For these reasons, the performance of the closed-loop system, defined here as the residual frequency noise, is directly related to high open loop gain and low noise, especially at the sensor.

Just how much gain can be practically implemented is restricted largely to the bandwidth of the control loop. The Nyquist criterion defines the gain and phase relations that must be met to insure closed-loop stability, which places limits on the slope of the open-loop gain curve near the unity-gain frequency [1, 29]. For instance, the unity-gain frequency is limited to ~ 20 kHz for piezo-mirror actuators like the one used in the experiments in Ch. 5 (to avoid mechanical resonances in the PZT stack). Stable phase requirements limit the slope and therefore the values of gain at the few kHz levels where acoustic-mechanical noise is commonly found (see Fig. 5.3, pg. 73 & Fig. 5.9, pg. 86 for Bode plots of open-loop transfer functions).

Although a particular actuator may well determine the feedback bandwidth, the sensor—or frequency discriminator—is a critical part of the ultimate performance limit [67]. The success of feedback systems in reducing the spectral width of CW NPRO lasers to sub-Hz levels is attributable to the fact that the residual frequency noise is reduced to near the detector quantum, or shot, noise [19].

This level of noise suppression is not necessary for driving an external enhancement cavity, however. An estimate of the required frequency stability of the drive laser can

be derived from the coupled oscillator (Adler) equations [7], but in general is a more complicated function of the residual laser spectral noise (for example Fig. 5.7 pg. 83). One plausible simplification is to model the laser noise as a distribution $g(\omega - \omega_0)$ of delta-function linewidth frequencies centered at the cavity resonance ω_0 . This model represents ideal resonance tracking but with some closed-loop error term that fluctuates with $g(\Delta\omega)$. The frequency response of a resonant cavity is given by [46],

$$E(\omega) = \frac{1}{\sqrt{2\pi}} \frac{1}{(\omega - \omega_0)i - \Gamma/2}, \quad (4.23)$$

where Γ is the FWHM of the cavity ($= \omega_0/Q \simeq \text{FSR}/b$). The total coupled power would then be described by this response multiplied by the distribution of frequencies,

$$P \propto \int |E(\omega)g(\omega)|^2 d\omega. \quad (4.24)$$

Even simple distributions $g(\omega)$ lead to complicated analytic expressions, but for distributions of width $\Delta\omega \ll \Gamma$, the coupled power must scale $\propto (1 + (\Delta\omega/2\Gamma)^2)^{-1}$. Therefore, the coupling efficiency is expected to reach near unity rather quickly as the residual noise is driven below the cavity bandwidth. In terms of the expected source cavity design, the cavity bandwidth will be a few kHz, which would conservatively require a few hundred Hz closed-loop residual noise bandwidth on the laser. This value is still a factor of a 1000 higher (easier) than what has been achieved for very stable lasers [83], and is a factor of 10 higher than the residual spectral widths of stabilized diode lasers with large, several MHz free-running noise [69].

4.3.3 Pound-Drever-Hall Frequency Discriminator

The essential idea of Pound-Drever-Hall (PDH) frequency stabilization is to use an external interferometer as a stable frequency reference, measure the frequency fluctuations of the laser against this reference, and then suppress them using feedback on either the cavity or laser. To accomplish this task, one needs a signed error signal representing the laser frequency fluctuation from that of the cavity resonance. The PDH technique supplies this discriminator signal by sampling the phase of the fields in the

cavity through optically mixing the reflected fields leaking out of the cavity with a reference phase given by promptly reflected fields of known modulated sidebands on the incident carrier.

There are two major benefits of using PDH over other stabilization techniques. First, it does not use another frequency reference source but rather the same optical cavity needed for gain enhancement. Second, by using reflected power, it has intrinsically high bandwidth, unlike transmission fringe-locking which has prohibitive time delays for high finesse cavities [19]. The external optical cavity can then perform two related but different functions—power amplification *and* frequency stabilization.

PDH Theory

The incident beam is first sent through an electro-optic phase modulator which puts rf sidebands, Ω , on the main optical frequency spectrum, ω . The modulated incident field has the form

$$\begin{aligned}\vec{E}_{\text{inc}} &= E_0 e^{i(\omega t + \beta \sin \Omega t)} \\ &\approx [J_0(\beta) + 2iJ_1(\beta) \sin \Omega t] e^{i\omega t} \\ &= E_0 [J_0(\beta) e^{i\omega t} + J_1(\beta) e^{i(\omega+\Omega)t} - J_1(\beta) e^{i(\omega-\Omega)t}],\end{aligned}\quad (4.25)$$

where β is the modulation depth in radians. For $\beta < 1$, most of the power is in the carrier and first-order sidebands only,

$$\begin{aligned}P_0 &\approx P_c + 2P_s, \\ \text{where } P_c &= J_0^2(\beta) P_0 \quad \text{and } P_s = J_1^2(\beta) P_0.\end{aligned}$$

The reflected field is then the incident field, Eq.4.25, multiplied by the complex reflection coefficient, $F(\omega)$, of the cavity,

$$\vec{E}_{\text{refl}} = E_0 [F(\omega)J_0(\beta) e^{i\omega t} + F(\omega + \Omega)J_1(\beta) e^{i(\omega+\Omega)t} - F(\omega - \Omega)J_1(\beta) e^{i(\omega-\Omega)t}].\quad (4.26)$$

For a FP cavity, the general form of the reflection coefficient is

$$F(\omega) = \frac{-r_1 + r_2(r_1^2 + t_1^2) \exp\left(i\frac{\omega}{\text{FSR}}\right)}{1 - r_1 r_2 \exp\left(i\frac{\omega}{\text{FSR}}\right)}, \quad (4.27)$$

where r_i and t_i are the amplitude reflection and transmission coefficients of the mirrors ($R_i = r_i^2$, $T_i = t_i^2$), and FSR is the free spectral range. Characteristic of a resonator, $F(\omega)$ sweeps out a circle on the complex plane, and for a matched cavity passes through the origin at resonance. A photodetector which measures reflected power, $|\vec{E}_{\text{refl}}|^2$, then contains several terms with frequencies at DC, Ω , and 2Ω . In the case that the modulation frequency is large compared to the cavity bandwidth, $F(\omega \pm \Omega) \approx -1$, the reflected power can be approximated near resonance to be

$$P_{\text{refl}} \approx (P_c |F(\omega)|^2 + 2P_s) - 4\sqrt{P_c P_s} \text{Im}[F(\omega)] \sin \Omega t + (2\Omega \text{ terms}). \quad (4.28)$$

The term we are interested in is the carrier-sideband mixing, proportional to Ω , which is extracted by demodulating the signal with $\sin \Omega t$ using a mixer. A typical error signal, ε , is plotted in Fig. 5.2 (pg. 72). The slope of the error signal near resonance is found by approximating the reflection coefficient for a high finesse cavity,

$$F \approx 2i \frac{\delta f}{\Delta \nu_{\text{cav}}},$$

where δf is the deviation of the laser frequency (in Hz) from resonance. The error signal can now be cast as a linear function of δf ,

$$\varepsilon = D\delta f, \text{ where } D \equiv -\frac{8\sqrt{P_c P_s}}{\Delta \nu_{\text{cav}}}. \quad (4.29)$$

This error signal now forms the basis of a nulled lock-in feedback system using the standard tools of linear control theory as outlined in Sect. 4.3.2.

One important trait of this PDH error signal is that the linear portion of the discriminator is valid only up to $\pm \Delta \nu_{\text{cav}}/2$. To insure a stable working point at resonance, the peak-to-peak residual frequency excursions should remain within these

linear bounds. A commonly used estimate for peak-to-peak frequency is to multiply the measured RMS frequency noise by 5. This scaling is consistent with the required frequency stability necessary for efficiently coupling power (Eq. 4.24). Conversely, if the residual laser noise is a significant fraction of the external cavity bandwidth, both the tracking performance and coupling efficiency are expected to degrade.

Another interesting result of Eq. 4.29 is that the intrinsic sensor gain scales inversely to the cavity bandwidth, or proportional to the cavity finesse. Therefore, it is not surprising that high performance, low noise, control loops prefer to use very high finesse cavities, as Eq. 4.22 predicts.

Chapter 5

Optical Cavity Experiments

Little work has been done previously to determine the performance and limitations of amplifying mode-locked laser pulses in high-finesse, passive optical cavities. The most relevant experiments in the literature report either modest cavity gain enhancements (gains ~ 100), or describe the use of optical cavities as interferometers for implementing laser frequency stabilization. In the first group are studies of optical pulse stacking, such as at the Stanford mid-IR FEL [17, 37, 74], or the use of external cavities to drive more efficient harmonic conversion with mode-locked lasers [59, 64]. On the other hand, frequency stabilization [72] by itself does not require cavity enhancement, but it is a prerequisite for stable, high gain operation. Therefore, demonstrating pulse-stacking with a high-power mode-locked laser is one of the most important experimental verifications necessary to validate the basic design of the laser-electron ring X-ray source.

There are two primary cavity tests that can be independently performed—achieving high gain, requiring longitudinal stability, and achieving a small waist, requiring transverse stability. Of course, both conditions must be met in a practical design, but separating these issues makes them experimentally more convenient. Near-concentric cavities with small waists are best studied using stable, single-mode CW lasers. CW lasers have the freedom of driving any axial mode for a cavity resonator, so adjusting the cavity waist given a set of mirror radii is straightforward. However, for pulse

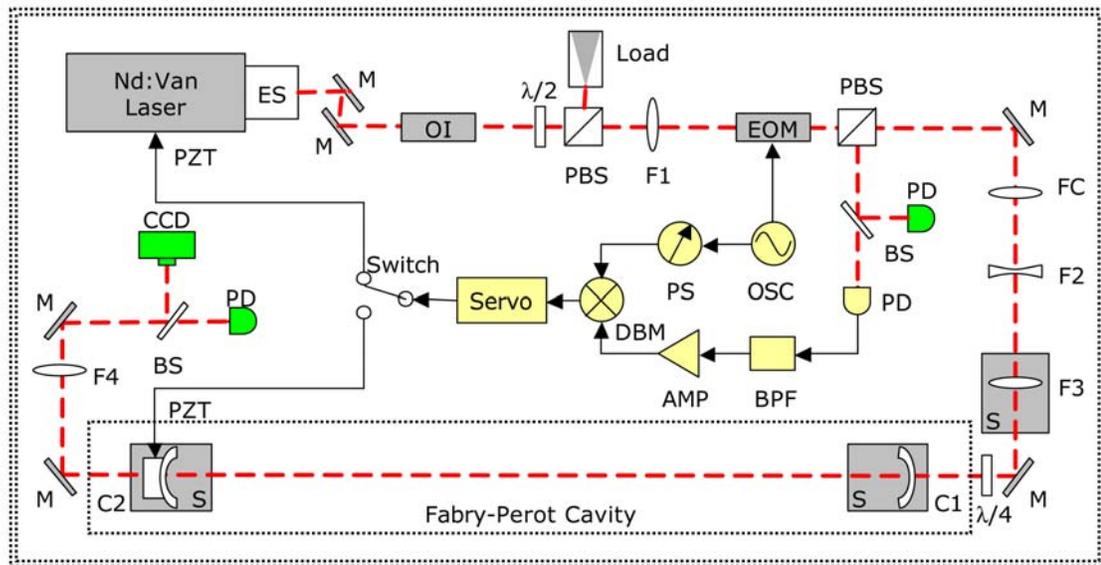


Figure 5.1: Experimental layout shown roughly to scale on 4x8 ft optical table. Beam propagates clockwise from laser to Fabry-Perot cavity (enclosed in Plexiglas cover). Key: ES = Electronic Shutter, M = Mirror, OI = Optical Isolator, $\lambda/2$ = Half-Wave Plate, PBS = Polarizing Beam Splitter, F = Lens, FC = Cylindrical Lens, EOM = Electro-Optic Modulator, $\lambda/4$ = Quarter-Wave Plate, S = Linear Stage, C = Cavity Mirror, BS = Beam Splitter, PD = Photodiode, OSC = Oscillator, PS = Phase Shifter, BPF = Band-Pass Filter, AMP = Amplifier, DBM = Double-Balanced Mixer, PZT = Piezo-Electric Mirror Mount.

stacking, the condition of matching the passive cavity length to that of the mode-locked laser cavity dictates that the mirror radii effectively determine the cavity waist. Because of the limited repetition-rate frequency tuning of the laser and limited choice of mirror curvatures commercially available, near-concentric, very small-waist geometries were not tested in the experiments reported here. The pulse-stacking amplification requirements, though, are identical for any stable cavity; larger-waist cavities are simply less sensitive to alignment and mode-matching errors but more sensitive to high-power handling since the spot sizes on the mirrors are reduced.

5.1 Experimental Apparatus

The optics layout, viewed schematically in Fig. 5.1, consists of the laser, the transport optics, and the FP cavity. The entire setup fits on one isolated optical table (*Newport* RS4000/I-2000) which is located in a class 10,000 cleanroom lab.

5.1.1 Mode-Locked Laser

The laser used in these experiments is a *High-Q IC-10000*—a 10 W, diode-pumped, Nd:YVO₄ CW passively (SESAM) mode-locked laser operating at a wavelength of 1064 nm. The laser is packaged in a compact, OEM-style housing and includes two factory modifications, one for timing synchronization and the other for lengthening the pulse.

The timing stabilization option allows a small tuning (< 100 kHz) of the nominal 79.33 MHz pulse repetition rate and is intended to phase-lock an external rf signal to the pulse train in order to reduce pulse-to-pulse timing jitter (≤ 0.5 ps). In practice, this system does little to improve the free-running frequency noise of the laser, and was quickly abandoned. However, one part of the system—the piezo-mirror mount in the laser cavity—eventually proved useful as part of the frequency stabilization feedback. The worst deficiency of the installed piezo is that it is not optimized for high-bandwidth applications since the timing synchronization electronics do not need much more than 1 kHz bandwidth. Even so, it is fortunate that the first serious mechanical resonance of the piezo measures out at nearly 20 kHz.

The other modification to the laser cavity is an added etalon filter. Since shorter pulses are more sensitive to dispersion effects (Sect. 4.2.3), the natural 7 ps pulse width was stretched to 25–30 ps. The etalon limits the available gain bandwidth, thereby reducing the number of axial modes forming the pulse envelope while still maintaining stable “transform limited” pulses. The predicted effects of dispersion for 30 ps pulses is quite small even for the highest finesse cavities tested (see Sect. 4.2.3).

5.1.2 Transport Optics

The optics train, following in sequence from the laser to the cavity (Fig. 5.1), starts with a pair of mirrors to set beam height (4 in.), a Faraday optical isolator (*OFR IO-8-YAG-VHP*) to prevent optical feedback to the laser, a half-wave plate/polarizing beam splitter attenuator to adjust power, and a lens f_1 to collimate the beam to a roughly 1 mm radius. Next, the beam passes through an electro-optic phase modulator (*EOPC EO-361P*) which adds primarily first-order frequency sidebands at a modulation rf frequency, Ω . The beam passes through another polarizing beam splitter, which combined with the quarter-wave plate at the cavity entrance acts as a directional coupler; reflections from the cavity rotate polarization and are then deflected at the beam splitter along the path toward the two photodetectors, one for monitoring power and the other used for frequency stabilization feedback. Lastly, the two mirrors and intervening lenses at the end of the table nominally align and transversely mode-match the beam to that of the cavity. The mirrors have *New Focus* picomotor adjustable mounts and the final lens f_3 is mounted on a linear stage to dynamically adjust the coupling waist size and position if necessary.

The optics at the output of the cavity are only for diagnostics—a photodiode for transmitted power measurements and a CCD camera to view mode patterns and assist in initial cavity alignment.

5.1.3 Fabry-Perot Cavity

The FP cavity is simply formed by two mirrors mounted directly on the optical table and locally covered by a Plexiglas hood to prevent air currents. Each of the mirrors has a computer-controlled picomotor mount fixed on top of a dc-servo linear translation stage (*PI M-126.DG*) capable of sub-nm resolution movement along the cavity geometric axis. In addition, the output mirror is attached to a tip/tilt piezo aligner (*Burleigh PZ-81*) providing both fine angular and cavity length adjustments ($2\ \mu\text{m}/1000\ \text{V}$).

All the cavities built and tested use commercially available mirrors: fused-silica substrate, dielectric coated, 1 in. diameter, with nominal 1.0 m radius of curvatures.

Table 5.1: Nominal characteristics of mirrors used for Fabry-Perot cavities.

Manufacturer & Type	Label	R	T	L
<i>CVI</i> PR series-95R	95	95 %	≤ 5 %	0.1–0.2 %
<i>CVI</i> PR series-99R	99.5	99.5 %	0.3–0.4 %	0.1–0.2 %
<i>Newport</i> SuperMirror-T	ST	99.97 %	200–300 ppm	50–100 ppm
<i>Newport</i> SuperMirror-F	SF	99.98 %	100 ppm	50–100 ppm

Table 5.2: Cavity mirror combinations and expected performance.

Config.	$M_1 \otimes M_2$	Finesse	$\Delta\nu_{\text{cav}}$	Coupling	Gain
I	95 \otimes 95	63	1.3 MHz	100%	20
II	95 \otimes 99.5	110	700 kHz	33%	66
III	99.5 \otimes 99.5	630	130 kHz	91%	140
IV	99.5 \otimes ST	1200	70 kHz	90%	500
V	ST \otimes ST	9700	8.2 kHz	95%	2400
VI	ST \otimes SF	13000	6.3 kHz	100%	4000
VII	SF \otimes SF	18000	4.4 kHz	82%	3300

Notes: M_1 is input mirror; listed coupling and gain are for a perfectly mode-matched and aligned incident beam.

The only differences between them are the various distributions of reflectivity, transmissivity, and losses. Table 5.1 summarizes the four types of mirrors investigated: the first two are partially reflective mirrors made with electron beam deposition (*CVI*), and the second two are ultra-low loss ‘SuperMirrors’ made by ion beam deposition (*Newport T* and *F* series). Table 5.2 is a progression of cavities tested with various mirror combinations in order of increasing finesse. The last configuration (VII) was not tested—the narrow cavity bandwidth makes the laser locking more difficult while the achievable gain of the system suffers because of the impedance mismatch.

5.1.4 Frequency Stabilization Feedback

The Pound-Drever-Hall technique generates a feedback signal by frequency modulating the optical beam, filtering and demodulating the reflected power signal from a fast photodiode, and then conditioning this error signal through a servo back to a mirror-mounted piezo actuator (see Sect. 4.3.3).

The modulation frequency, Ω , should be greater than the bandwidth of both the cavity and of the feedback control loop, but less than some sampling factor of the pulsed repetition rate of the laser. The optimal range is then likely between 1–10 MHz; 1.7 MHz was chosen for these experiments mainly because it was near the resonance of a step-up transformer required to drive the non-resonant electro-optic modulator (EOM). With the available rf amplifier, a modulation depth of 0.6 radians is typical for these experiments, resulting in a reduction of main carrier power by approximately 20%. The modulation depth can be tailored lower, however, to increase overall coupling efficiency; some later experiments used $\beta = 0.2$ rad. ($\approx 2\%$ sideband power) without noticeable effect on performance.

The reflected power from the cavity is directed on to a fast photodetector (*Thorlabs* 410) where the intensity is adjustable using ND filters. The typical power on the detector is around 1 mW. A 1 MHz high-pass filter in series with a 1.9 MHz coax low-pass filter removes the DC component as well as the 79.33 MHz pulsed component. A broadband 20 db rf amplifier (*Mini-Circuits* ZFL-500) is placed in-line immediately after the filters to boost the signal for easy viewing on a scope. This signal then feeds a double-balanced mixer (*Mini-Circuits* ZAD-6BR) which is demodulated with a properly phased sine wave at Ω to recover the DC component. The signal is subsequently low-pass filtered with an elliptic filter notched to remove any mixed 2Ω term. The remaining signal is now called the error signal, ε , which reproduces the classic PDH form as seen in Fig. 5.2. It is continuously monitored on a scope and FFT analyser (*Stanford Research Systems* SR785) for evaluating the closed-loop frequency noise suppression.

The error signal is the input to a servo compensator which conditions the signal for stable closed-loop feedback to a mirror-mounted piezo actuator. The working servo for early, large bandwidth cavity configurations uses a single pole, 10 Hz roll-off

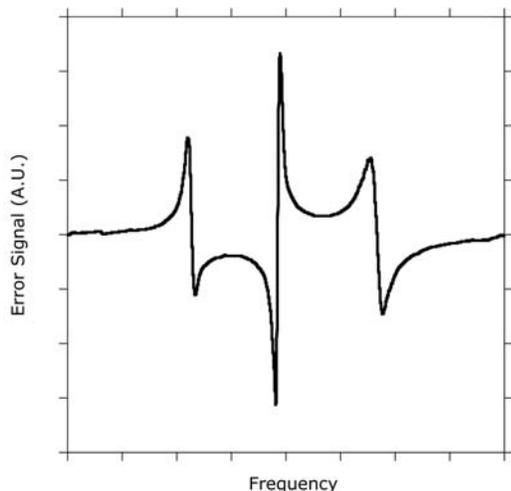


Figure 5.2: Example of a PDH error signal seen by applying a linear frequency sweep of the FP cavity through a resonance with a free-running laser. The sidebands ($\beta = 0.6$) are resonant 1.7 MHz away from the central carrier where the slope is linear over the cavity bandwidth (shown here $\Delta\nu_{\text{cav}} = 126$ kHz). In practice, the measured timing between resonances, and even resonant widths, vary sweep-to-sweep due to the RMS frequency noise of the laser superimposed on the cavity ramp rate.

low-pass filter, where the unity gain frequency is set just below the first major piezo resonance. The optimal settings were found empirically by using a *Stanford Research Systems* SR560 low-noise pre-amp that has adjustable gain and filter settings. To obtain a large gain (~ 80 dB) at low frequencies, driving the piezo capacitance with the $600\ \Omega$ output impedance port works best by providing a steeper gain roll-off toward the unity gain frequency while still ensuring plenty of phase margin (Fig. 5.3). The narrow-band cavity configurations required a better performing servo which is tailored to the noise spectrum of the laser (see Sect. 5.3.2).

There are two possible actuators in the system—the tip/tilt piezo on the cavity output mirror, and the factory installed piezo mirror in the laser. For several reasons, the laser piezo was used almost exclusively. First, although it’s not as stiff as the high-voltage *Burleigh* piezo, it drives a much lighter mirror. The measured resonances on the FP cavity mirror piezo start at 5.5 kHz, while the laser piezo has no dramatic resonances until approximately 20 kHz. The higher bandwidth available in the laser piezo became a critical factor in locking the laser to the narrow-band cavities. Secondly, the laser is the dominant source of frequency noise, so making the intrinsically stable passive cavity track a noisy source is not as natural as suppressing the noise in the laser to that of the cavity.

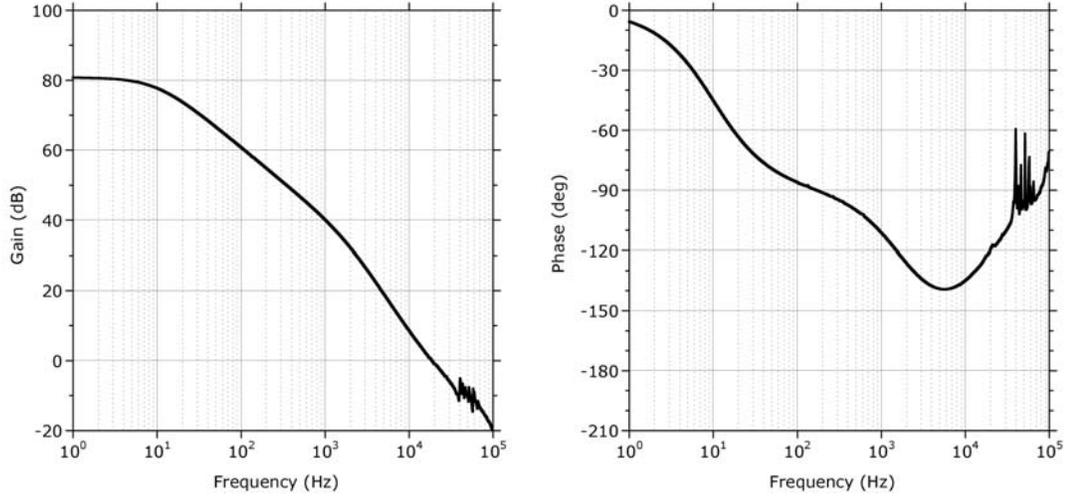


Figure 5.3: Measured open loop frequency response of the servo used to initially lock to broad-band cavities (I–IV, Table 5.2). The servo is a 10 Hz low-pass filter combined with the additional low-pass filter effect of driving a piezo capacitance in series with the SR560 600 Ω output impedance.

5.2 Experimental Procedure

5.2.1 Cavity Mode-Matching

With a fixed cavity length, $L = 1.89$ m, and chosen mirror radii, $R = 1.0$ m, the TEM_{00} eigenmode of the cavity is fully determined. For our symmetric FP cavities, the axial waist location, z_0 , is centered between the mirrors and has a size (Eq. 4.7),

$$w_0 = \sqrt{\frac{\lambda L}{2\pi} \left(2\frac{R}{L} - 1\right)^{1/2}} \approx 280 \mu\text{m}. \quad (5.1)$$

The laser input must match this waist size and position to best couple power. For a real laser, the beam is not ideally a pure Gaussian but a superposition of Laguerre-Gauss resonator modes and is characterized by the M^2 value ($M^2 \geq 1$) where

$$w_0 = W_0/M, \quad (5.2)$$

Table 5.3: Measure of laser mode-matching fit to cavity.

Parameter	\hat{x}^\dagger	\hat{y}^\dagger	Error	Cavity
z_0^* (mm)	958	923	± 1	944.8
W_0 (mm)	0.304	0.301	± 0.002	0.2781
z_R (mm)	242	229	± 4	228.4
M^2	1.13	1.16		1.00

*measured from input mirror.

†axes defined relative to ellipticity of beam.

such that the measured waist W_0 is M times larger than the embedded Gaussian in both \hat{x} and \hat{y} [48]. The appropriate parameter to match is then the Rayleigh range, z_R , related to the focusing depth and common to all the modes,

$$z_R = \frac{\pi w_0^2}{\lambda} = \frac{\pi W_0^2}{M^2 \lambda} \quad (5.3)$$

Therefore, the general strategy is to characterize the laser, including its M^2 , and place the optics necessary to focus the waist at the proper location and focus depth.

Transverse Mode-Matching Results

The laser was first corrected for astigmatism ($z_{0x} \neq z_{0y}$) and slight asymmetry ($w_{0x} \neq w_{0y}$) by adding a cylindrical lens, f_c , upstream of the mode-matching lenses f_2 and f_3 (Fig. 5.1, pg. 67). The two mode-matching lenses formed a telescope such that moving their relative separation changed the focus depth while moving them together changed the waist location.

After placing the optics in their calculated positions—including an input mirror which acts a slight defocusing lens—the cavity waist was measured *in situ* and positioned as close as possible to the geometric cavity center. Using a simple “4-cuts” method [48] to fit the waist parameters, several small iterations in lens placement yielded the final matching values shown in Table 5.3.

Small mode-matching errors of the input beam excite some of the next radial higher-order cavity modes. However, near resonance of the fundamental mode, the

cavity effectively filters *all* higher mode content. Practically, this means that power into any but the fundamental mode is promptly reflected at the input mirror and simply reduces coupling efficiency (see Sect. 4.2.1). The mode-matching error parameters are measured waist size, w'_0 , and axial waist position displacement, Δz_0 . The effective TEM₀₀ waist, w'_0 , from Table 5.3 is given by the measured $W_0/\sqrt{\langle M^2 \rangle}$. Calculating these errors using the relevant expansion coefficients (Table 4.1, pg. 51), the expected power losses are dominated simply by the M^2 quality of the beam. The average M^2 is taken to be 1.15 for a baseline TEM₀₀ coupling efficiency of 87%.

5.2.2 Cavity Alignment

Once the mode-matching optics were optimized, alignment of the input beam to the cavity optical axis was performed in the following sequence of steps (and generally repeated for each set of new cavity mirrors).

1. Align the laser through the geometrical axis of the cavity by placing apertures¹ in the cavity mirror mounts. The CCD camera at the output end is aligned to this cavity reference.
2. Set the input mirror, M_1 , and retro-reflect back to the optical isolator by removing the $\lambda/4$ waveplate (which when replaced has negligible effect on beam steering).
3. Set the output mirror, M_2 , and place an aperture in the cavity to align the first reflection off M_2 back along cavity axis. After removing the aperture, some weak resonances of higher-order modes is usually observed at the output by which the TEM₀₀ mode can be found.
4. Slowly scan (\sim Hz) cavity frequency using a mirror piezo, and tune-up the TEM₀₀ power by iteratively using the two cavity mirrors to adjust the cavity axis tilt (differential mode) or offset (common mode) for \hat{x} and \hat{y} .

¹6 mm thick aluminum 1.000 ± 0.001 in. optic ‘blanks’ with a 0.8 mm centered hole ($\pm 25 \mu\text{m}$)

The transmitted mode patterns on the CCD now represent the cavity eigenmodes only. One way to minimize the alignment error is to realize that small errors predominantly couple to the next mode (TEM_{01}) which can be minimized relative to the fundamental by watching the mode strengths on the CCD or comparing photodetector signals.

The errors associated with misalignment at this point are dominated by the uncertainty ($\sim 10\%$) of the peak transmission maximum due to amplitude fluctuations of the transmitted power caused by the laser free-running frequency noise. Stabilizing the laser and locking to the fundamental cavity mode would then allow a better optimization of coupling.

Setting the Absolute Cavity Length

During alignment, the overall cavity length was scanned near the absolute length required to match the laser repetition rate. For gross tuning, the timing stabilization electronics in the laser controller was monitored, but not enabled, so that there was a continuous display of the laser cavity frequency compared to a set 79.333 MHz rf reference source. Warm-up drifts of 2–3 kHz were typical but long-term stability was generally ~ 100 Hz/Hr.

By adjusting the linear stage on one cavity mirror and looking at the transmitted power signal, it was fairly straightforward to maximize the peak, which also minimizes the width, of the photodiode signal. The sensitivity of setting the absolute length correctly is dependent on cavity bandwidth—the broad-band cavities are tolerant of $\lesssim 10\mu\text{m}$ errors while the narrow-band cavities performed best within a few axial modes of the peak. Fortunately, the narrow-band cavities exhibit additional features that helped locate the optimal axial mode (see Sect.5.3.1).

5.2.3 Measuring Cavity Parameters

The moderate finesse ($\sim 10^4$) cavities used in these experiments must be characterized *in situ* since the only mirror attribute independently measurable is the small transmission, T , leaving the reflectivity $R \approx 1$ with small but difficult to measure losses. Two techniques are discussed: the first uses simple power measurements and

proved quite robust at characterizing both the cavity and input coupling; the second, based on the dynamics of frequency swept cavities, yielded less consistent results but nonetheless proved useful in making a quick determination of mirror losses by direct estimation of cavity bandwidth in one “snapshot” of the transmitted power.

Reflected and Transmitted Power

The most straightforward measure of cavity gain is a calibrated set of power measurements of the reflected and transmitted photodiode signals. Knowing the total incident power and the transmission of each mirror, one can algebraically solve for the two unknowns in the system: the complete losses in the cavity, and the coupling ratio—that fraction of the incident beam which, when expanded in the eigenfunctions of the cavity [6], can couple to the fundamental TEM₀₀ cavity mode. Just how much of that power actually couples into the cavity is then dependent on the impedance match (see Sect. 4.1.2).

Expanding on a short note by Smith and Shernoff [73], one can derive the relevant expressions necessary to solve for the cavity bounce number, b , and the TEM₀₀ coupling factor, c_0 , in terms of the intracavity fields through purely external measurements: Let the power matched to the proper cavity eigenmode of an incident beam P_{inc} be defined as

$$P_0 = c_0 P_{\text{inc}} \quad (5.4)$$

where *a priori* only the incident power is measured and the coupling coefficient c_0 is unknown. The measured transmitted power P_{tran} is directly related to the intracavity circulating power

$$\begin{aligned} P_{\text{tran}} &= P_{\text{circ}} \mathsf{T}_2 \\ &= 4c_0 P_{\text{inc}} b^2 \mathsf{T}_1 \mathsf{T}_2 \end{aligned} \quad (5.5)$$

using Eq. 4.3 for the cavity gain. The measured reflected power is a combination of the prompt mismatched beam and the net cavity reflection given by the fields described

in Eq. 4.4,

$$\begin{aligned} P_{\text{refl}} &= (P_{\text{inc}} - P_0) + P_0 \left(\frac{\vec{E}_{\text{refl}}}{\vec{E}_{\text{inc}}} \right)^2 \\ &= (1 - c_0) P_{\text{inc}} + c_0 P_{\text{inc}} (1 - 2\mathbb{T}_1 b)^2. \end{aligned} \quad (5.6)$$

Use Eq. 5.5 and Eq. 5.6 to solve for b and c_0 ,

$$b = \frac{P_{\text{tran}}}{P_{\text{inc}} \mathbb{T}_2 + P_{\text{tran}} \mathbb{T}_1 - P_{\text{refl}} \mathbb{T}_2}, \quad (5.7)$$

$$c_0 = \frac{(P_{\text{inc}} \mathbb{T}_2 + P_{\text{tran}} \mathbb{T}_1 - P_{\text{refl}} \mathbb{T}_2)^2}{4P_{\text{tran}} \mathbb{T}_1 \mathbb{T}_2 P_{\text{inc}}}. \quad (5.8)$$

Eq. 5.7 directly yields the cavity finesse, $2\pi b$, and upon rearranging, a simple expression for the total losses in the cavity and for the aggregate scattering and absorptive losses for the mirrors,

$$\text{total cavity loss} \approx \frac{1}{b} = \mathbb{T}_1 + \left(\frac{P_{\text{inc}} - P_{\text{refl}}}{P_{\text{tran}}} \right) \mathbb{T}_2, \quad (5.9)$$

$$\text{combined mirror losses} \approx \mathbb{L}_1 + \mathbb{L}_2 = \left(\frac{P_{\text{inc}} - P_{\text{refl}} - P_{\text{tran}}}{P_{\text{tran}}} \right) \mathbb{T}_2. \quad (5.10)$$

This technique can be scaled to folded cavities where there are N round-trips by the circulating pulse per input pulse. Without redoing the algebra, one can use the above relations as is and then divide the calculated b by N and multiply the calculated coupling c_0 by N to obtain the proper results.

Note that all measurements are ratios of powers, so the absolute accuracy of the power meter used to calibrate the photodiodes is less important than the relative accuracy over the dynamic range. In these experiments, all photodetectors were calibrated against a *Coherent* Ultima Labmaster power meter (with a LM-2 IR semiconductor sensor and fixed 1000:1 attenuator) which has an overall accuracy of 5% but is linear to better than 1%. Also, despite the various coupling unknowns, the intrinsic finesse of the cavity can be measured with good resolution and is generally repeatable over a wide range of measured couplings, *e.g.* when the laser is locked to,

vs. simply free-running through, cavity resonance.

Note on Coupling The factors that affect coupling c_0 can be divided into two timescales: the first is essentially static and includes the mode quality of the laser, the modulation depth of the sidebands for the PDH signal, and the alignment and mode-matching of the transport optics to the cavity; the second is dynamic and results from the phase and frequency stability of the laser over the timescales near the fill-time of the cavity. Even when the laser is tracking the cavity central frequency, the coupling dynamics becomes quite noticeable when the laser RMS integrated noise is comparable to the cavity bandwidth.

Measuring Swept Pulse Dynamics

Several papers discuss the dynamics of Fabry-Perot cavities with respect to linear sweeps of either the input laser frequency, or similarly, the cavity length [53,66]. There are three regions of behavior one can qualitatively describe for different sweep rates. For “slow” sweeps, where the time to sweep across resonance is much longer than the cavity fill-time, the cavity can always be treated in steady-state and leads to the usual resonance peaks on the measured transmitted power. At the other extreme, for “very fast” sweeps, the energy coupled into the cavity acts like an impulse; the transmitted power then appears as an exponential decay with a ring-down time characteristic of the cavity bandwidth. For sweep times on the order of the cavity fill-time, the linear frequency sweep adds a non-linear phase dependence in the circulating fields which leads to additional peaks at stationary phase points, *i.e.* the appearance of significant “ringing” on the transmitted and reflected power.

The main features of the ringing are dependent on two parameters, the cavity bandwidth and the sweep speed. By knowing the frequency ramp rate, the cavity bandwidth can be measured directly. Unfortunately, for these experiments, the free-running RMS fluctuations of the laser produced significant variations in the cavity sweep speeds at the values where these effects are measurable. One technique proposed by Poirson *et al* [66] circumvents knowing the actual frequency sweep rate by taking a snapshot of the transmitted pulse and fitting a relation between the ratio of

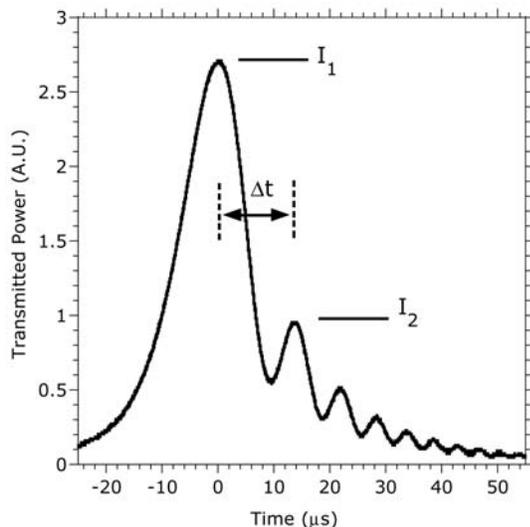


Figure 5.4: Example of ringing on the transmitted photodiode signal when sweeping cavity frequency. The ratio between the first two intensity peaks (I_1/I_2) is predicted to be a linear function of the time delay (Δt) between them. This snapshot is taken for a 10 Hz triangular sweep on the cavity piezo for a 12 kHz bandwidth cavity; it is near the case where $\Delta t \approx \tau (= \Delta\omega_{\text{cav}}^{-1})$.

the first two intensity peaks compared to the time delay between them. The predicted linear fit takes the simple form

$$\Delta\omega_{\text{cav}}\Delta t = \frac{1}{2} \left(\frac{I_1}{I_2} + 2 - e \right) \quad (5.11)$$

where $\Delta\omega_{\text{cav}} = 2\pi\Delta\nu_{\text{cav}}$. At the particular intensity ratio $I_1/I_2 = e$, the measured time delay Δt is exactly the characteristic decay time of the cavity, $\tau = \Delta\omega_{\text{cav}}^{-1}$ (see Fig. 5.4). A study of such measurements is discussed in Sect. 5.3.3.

5.3 Experimental Results and Discussion

The first series of cavity experiments focused on characterizing the intrinsic cavity gains. Since the power measurements in Sect. 5.2.3 separate cavity losses from coupling issues, all the cavity configurations could be reasonably measured with the free-running laser, even when the circulating fields never reach steady-state.

The problem then turned to determining the free-running noise of the laser. An order-of-magnitude estimate of the noise was first found by measuring the beat frequency of the mode-locked laser against a relatively more stable CW NPRO Nd:YAG reference (*Lightwave Electronics* model 125). The measured Nd:YVO₄ wavelength

(1064.38 nm) is just at the maximum tuning range of the NPRO (1064.37 nm) to ensure there is enough power to see a strong beat frequency between 0–80 MHz on an rf spectrum analyzer. At an analyzer sweep speed of 50 ms, a forest of very narrow modes covered approximately a 1 MHz span. Although this value is a few hundred times greater than the NPRO natural frequency stability, it is well within the dynamic range of the piezo-mirror frequency stabilization feedback for the mode-locked laser.

The cavity experiments were then essentially repeated, where the focus was now to acquire “lock” between the cavity and laser by suitable feedback design. The various cavity configurations are qualitatively grouped by cavity bandwidth in the following summary of results.

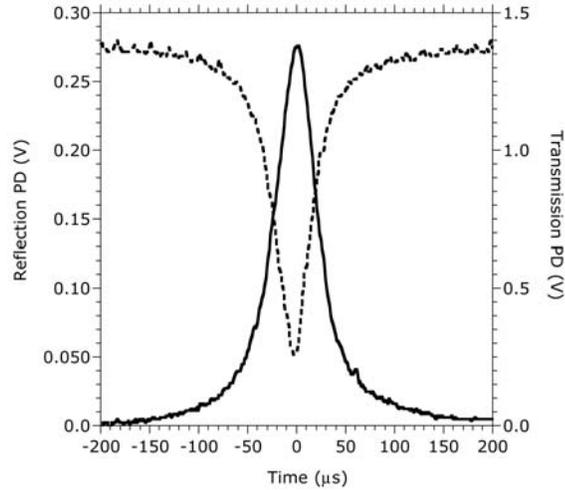
5.3.1 Broad-band Cavity Performance

The first set of cavity configurations (I–IV, Table 5.2) all display the same qualitative behavior. For large cavity bandwidths ($\gtrsim 70$ kHz), the short fill-time of the cavity allows the free-running laser to reach near steady-state values during slow frequency sweeps of the cavity; using the reflected and transmitted photodetector signals, both the coupling and losses of the cavity are easily measured (Fig. 5.5). A large cavity bandwidth also makes the transmission peak less sensitive to matching the exact longitudinal mode.

The basic servo (pg. 73) generally worked well enough to lock the laser to the cavity for many seconds, up to a few minutes, and was limited primarily by slow frequency drifts outside the dynamic range of the actuator drive voltage. No attempt was made to optimize the system long-term stability for these experiments, although it is not technically difficult to do so.

The quality of lock was optimal over a narrow range of gain settings and was a function of two competing aims: to minimize the integrated low frequency spectral noise (< 20 kHz), and to minimize the effects of the piezo actuator resonances (> 20 kHz). Reducing the large amplitude, low frequency noise produces more uniform transmitted and reflected signals by reducing the laser’s overall RMS frequency

Figure 5.5: Reflected and transmitted photodiode signals for a broad-band cavity using a free-running laser scanning through resonance. This cavity was formed from a 99.5×99.5 pair where $T_1 = 0.438\%$, $T_2 = 0.306\%$, and the calculated parameters are $b = 94.7$ and $c_0 = 0.84$. The added losses are therefore 0.31% , the cavity bandwidth $\Delta\nu_{\text{cav}} = 133\text{ kHz}$, and the gain ≈ 160 . The coupling is near the expected maximum considering the M^2 of the beam (Sect. 5.2.1).



excursions. However, at some servo gain the piezo mechanical resonances “turn on” and begin to frequency modulate the laser. As long as these resonances fall within the cavity bandwidth, some fraction of that power still enters the cavity and the reduction in coupling is minimized. In fact, the best stability often occurred while purposely “overdriving” the system—the improved low frequency noise allowed the laser to track the cavity more effectively while the piezo resonances acted as sideband modulations that still coupled power well inside the cavity bandwidth (Fig. 5.6).

Once the laser is stabilized well enough to lock to a cavity, the spectral noise of the closed-loop error signal, ε , measures the level of residual frequency noise (see Sect. 4.3.2). Also, by knowing the servo transfer function, this data gives a good estimate of the free-running laser noise spectrum (where the noise introduced by the discriminator is still considered small compared to the remaining laser noise). The overall magnitude of the free-running noise spectrum effectively determines the required servo bandwidth necessary to suppress the residual noise to a target RMS level. The RMS laser frequency noise as a function of Fourier frequency is calculated from the amplitude noise spectrum $S_{\text{laser}}(f)$ as

$$\bar{f}_{\text{laser}} = \sqrt{\int_f^{f_u} |S_{\text{laser}}(f)|^2 df}, \quad (5.12)$$

Figure 5.6: Reflected and transmitted photodiode signals for a laser locked to a 126 kHz BW cavity. The reflected signal shows the remaining frequency fluctuations while the transmission signal is filtered by the cavity. The servo gain was increased to the point of just starting a sharp 38 kHz piezo resonance. The measured coupling ratio dropped 10–15% from the predicted case, but the residual RMS noise improved (below).

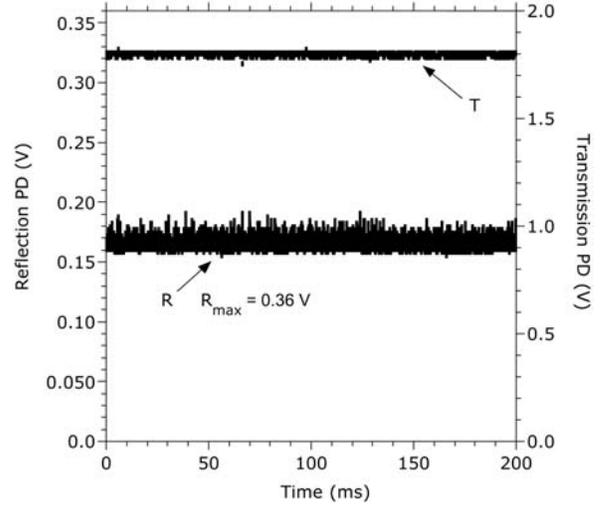
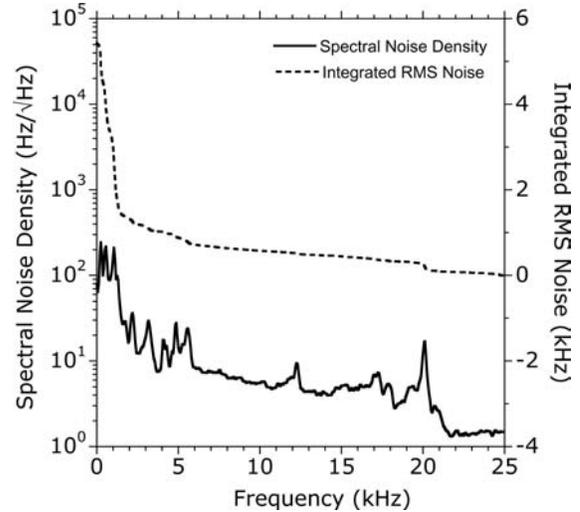


Figure 5.7: Spectral noise, $S_{\text{laser}}(f)$, of the closed-loop error signal for the 126 kHz BW cavity shown in Fig. 5.6. The 20 kHz piezo resonance could sometimes remain suppressed at gains where a strong resonance near 38 kHz dominated (not shown). Note that integrated noise is on a linear scale and that the major contributions over this range come in at low frequencies.



which is plotted on the right axis in Fig. 5.7 (where f_u is the upper frequency limit of the measurement). The value integrated over all frequencies defines the residual laser noise which is ideally suppressed to less than the external cavity bandwidth for efficient coupling and feedback stability.

5.3.2 Optimizing the Feedback Servo

As the bandwidth of the cavity decreases, the tracking requirement subsequently demands a lower integrated residual noise to keep the laser locked to the cavity. Ideally,

the gain would simply be increased, and to keep the same phase margin for stability, the unity-gain frequency would then increase too, which means increasing the overall feedback bandwidth. In these experiments, however, the onset of piezo mechanical resonances determined the maximum control bandwidth. Given this constraint, the servo needs to better match the actual noise spectrum of the laser in order to obtain low residual noise without exciting high-frequency resonances. Especially for the high-bandwidth cavity configurations, the induced piezo resonances would fall outside the cavity bandwidth and at best, reduce coupling power, and at worst, prevent locking altogether.

The basic strategy for optimizing the servo derived from studying the noise profile of the laser (Fig. 5.7). The bulk of the noise occurs at low kHz frequencies, consistent with the hypothesis that the noise sources are primarily mechanical resonances that couple to the optical beam. In particular, the goal was to minimize the three major resonances (at 250 Hz, 530 Hz, and 1060 Hz) that contributed most to the integrated noise.

First, an attempt was made to locate the source of the resonances in order to fix them passively. Measurement of the spectral noise of the diode-pump currents (2 heads, 3 V/60 A) revealed some harmonics of the 60 Hz line frequency but no correlations with the noise resonances. Next, a spare piezo mount was purposely driven at those fixed resonant frequencies and touched to the optical supports in the transport line while the cavity was locked. None of the external components seemed to be responsible for these noise resonances. The most sensitive piezo position turned out to be on the laser housing, near the output, but no attempt was made to investigate the optical supports in the laser cavity itself. The water cooling circuit may likely contribute some noise as well, but reducing the flow had deleterious effects on the stability and performance of the laser and so was left unchanged.

The task then was to build a compensator better tailored to the laser noise (Fig. 5.8). Since system stability was secondary to achieving lock, the servo was designed to have the steepest possible gain slope from the unity gain frequency given basic control theory stability requirements (max. slope ≈ 12 db/octave). Next, several parallel “booster” paths were placed within the servo circuit that added extra

resonant gain stages at exactly those frequencies where there is noise. By adjusting the gain and relative Q 's of these circuits, the phase excursions could be kept within stability bounds.

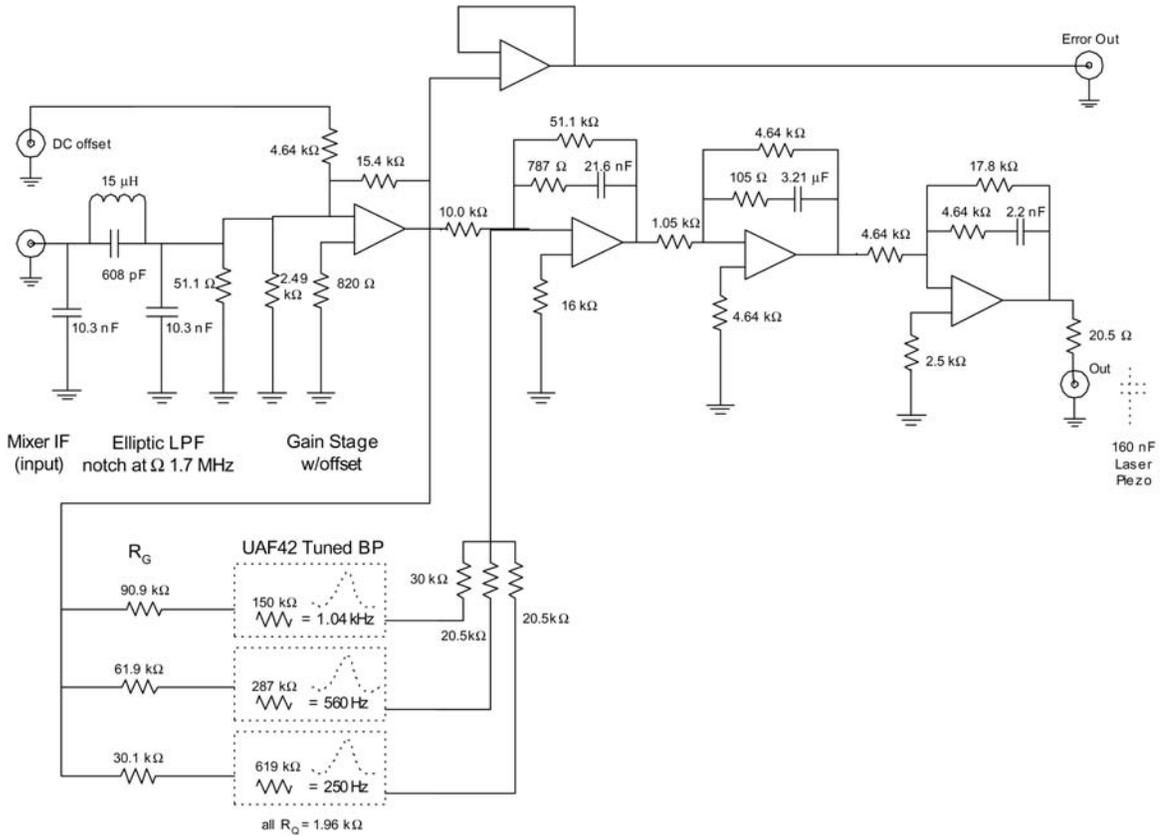


Figure 5.8: Servo electronics sketch for narrow-band cavity locking. The PDH error signal from the IF mixer port is filtered and amplified before feeding three parallel band-pass circuits tuned to the dominant noise spikes. The summed signal then is conditioned by a sequence of cascaded filters (using inverting LF356 opamps) in which poles and zeroes are added in pairs to optimize the shape of the overall response. Since the laser piezo is unipolar, the ± 15 V dynamic range of the opamps is biased by adding two 9 V batteries in series at the output (not shown).

Lastly, for narrow-band cavities, the calculated servo transfer function has to incorporate the effects of the error detection itself. The error term derived from the optical cavity behaves like a ‘frequency’ discriminator for excursions within the

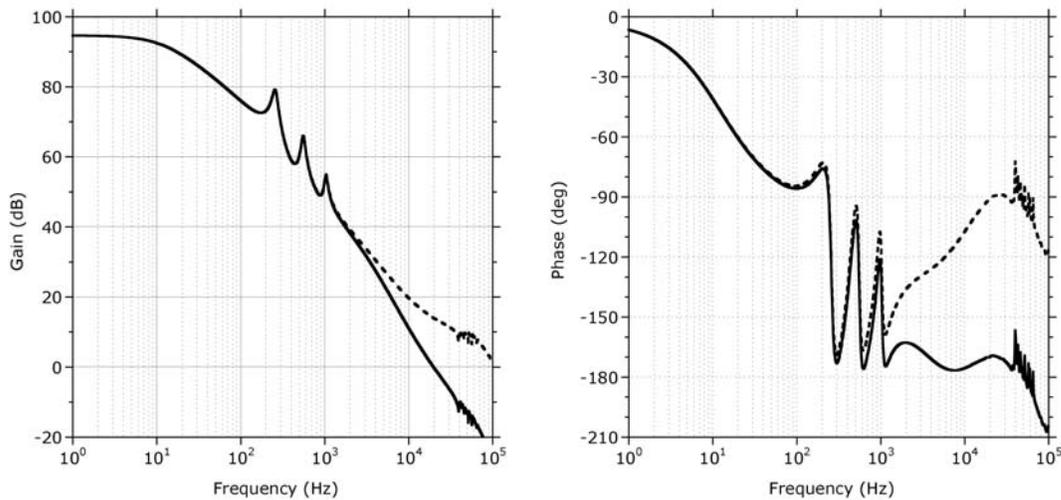


Figure 5.9: Open loop frequency response of the servo used to lock to narrow-band cavities (V & VI, Table 5.2). The dashed line is the measured response while the solid line predicts the total response by including the cavity low-pass filter effect on the discriminator signal (shown here optimized for an 8 kHz cavity bandwidth).

cavity half-bandwidth, and a ‘phase’ discriminator for excursions outside the cavity half-bandwidth [19,60]. In practical terms, the cavity behaves as a low pass filter with a roll-off at $\Delta\nu_{\text{cav}}/2$ which must be included in the overall loop frequency response when evaluating the total system transfer function (Fig. 5.9).

5.3.3 Narrow-band Cavity Performance

The tested narrow-band cavities (6–12 kHz) share some common qualitative features distinct from the broad-band cavities. The free-running laser could very rarely couple more than a fraction of its power to these cavities, even when matched; the long fill-times for these cavities prevented the input coupling from ever reaching steady-state values. Nevertheless, as long as the reflected and transmitted power had some short dwell time, the calculations for cavity losses and coupling were consistent (Fig. 5.10).

Early locking success occurred only for “detuned” cavities where the axial mode was purposely shifted from the peak by a few wavelengths. Besides reducing the transmitted power, the timing slip of the pulse envelope increases the effective bandwidth

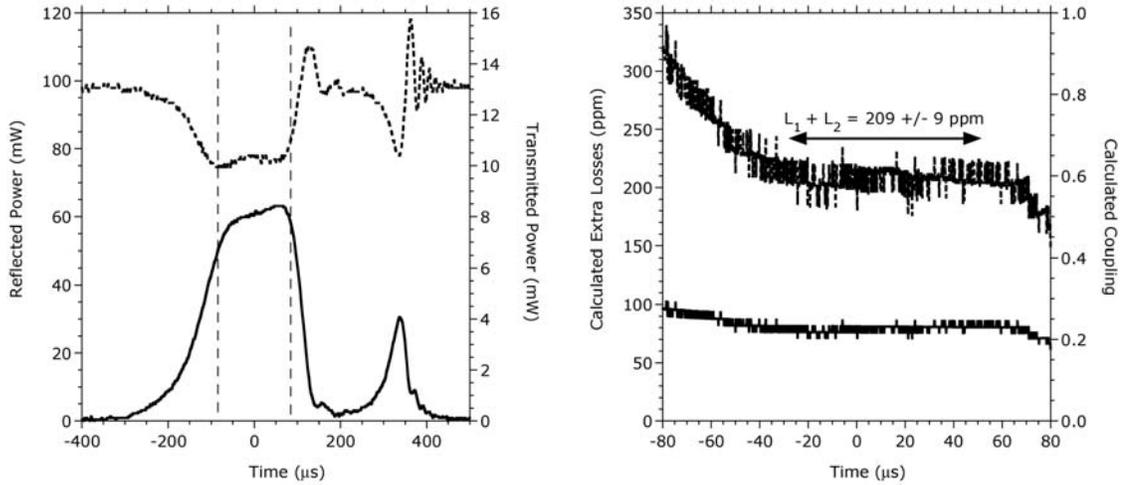


Figure 5.10: Calibrated reflected and transmitted power signals for a free-running laser passing through resonance in a narrow-band cavity. This cavity was formed from a $ST \otimes SF$ pair where $T_1 = 216$ ppm and $T_2 = 130$ ppm. The data within the dashed vertical bars are used to calculate the cavity parameters point-by-point in order to obtain a good mean value in a dwell region (shown on the right). These results give $b = 1810$ and $c_0 = 0.23$. The added losses are therefore 209 ppm, the cavity bandwidth $\Delta\nu_{\text{cav}} = 7.0$ kHz, and the gain ≈ 2800 . Note the extra transmission peak later in the waveform is another TEM_{00} resonance caused by the sometimes fast RMS frequency drifts in the laser imposed on the slow cavity frequency sweep.

of the cavity; both the fill-time and the number of pulses contributing to the steady-state circulating fields are decreased as the envelope slip-rate rises (see Sect. 4.2.3). As the axial mode was cycled closer to the peak, the lock performance degraded. Fig. 5.11 and shows the behavior of the locked cavity near the peak axial mode for a 10 kHz BW cavity along with the resultant spectral noise (Fig. 5.12). The features of the residual noise spectrum are also more complicated and show the effects of both the servo and piezo since the phase margin has been degraded.

The increased difficulty in acquiring lock as the axial modes were scanned towards the peak mode is directly related to the tracking requirement that the residual RMS laser noise be some small fraction of the cavity bandwidth. For the cavities tested here, the threshold bandwidth for more-or-less stable operation was approximately 20 kHz, or about 4–5 times the residual RMS integrated noise (see Fig. 4.6, pg. 55

Figure 5.11: 10.4 kHz cavity BW lock near peak axial mode. Both the stability and average coupling degrade as the axial modes are scanned closer to matching the laser repetition rate. This behavior is expected when the effective bandwidth of the cavity decreases while the residual error stays fixed.

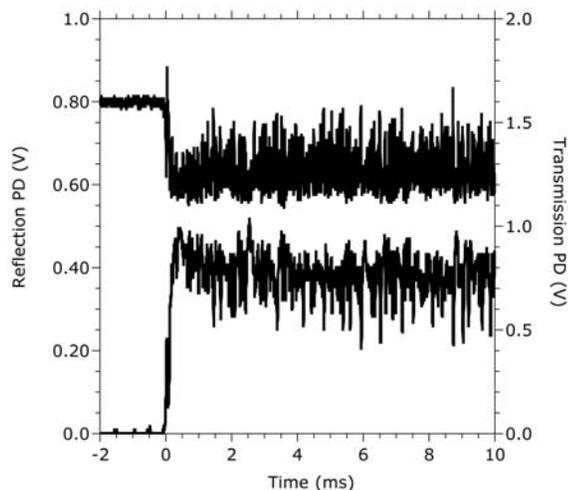
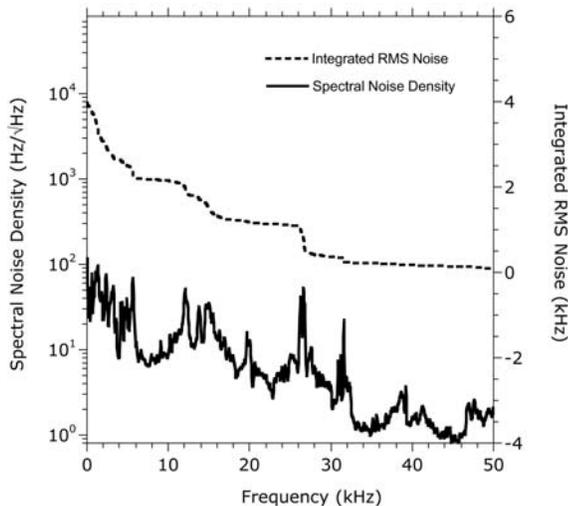


Figure 5.12: Spectral noise of the closed-loop error signal for a 10.4 kHz BW cavity. The low frequency laser noise resonances are now damped but there are additional features from the servo (10–20 kHz range) as well as piezo resonances (20–30 kHz range) that make the residual noise more complicated to optimize. The lowest achievable residual RMS noise was approximately 3–4 kHz.



for typical bandwidth vs. axial mode dependence). This scaling is consistent with the expected peak-to-peak bounds necessary to capture the residual error within the linear portion of the discriminator. With the broad-band cavities, reducing the low frequency noise at the expense of exciting piezo resonances was tolerated by the cavity. The narrow-band cavities, however, were more sensitive to any transients or lock instabilities. The mechanism responsible for poor locking performance is likely frequency excursions outside the linear range of the PDH signal, exacerbated by the dynamics of the cavity response; the error signal qualitatively changes when the residual RMS noise causes the frequency to sweep in and out of resonance at timescales near the cavity fill-time, the effects of which can be seen in Fig. 5.13. Although the

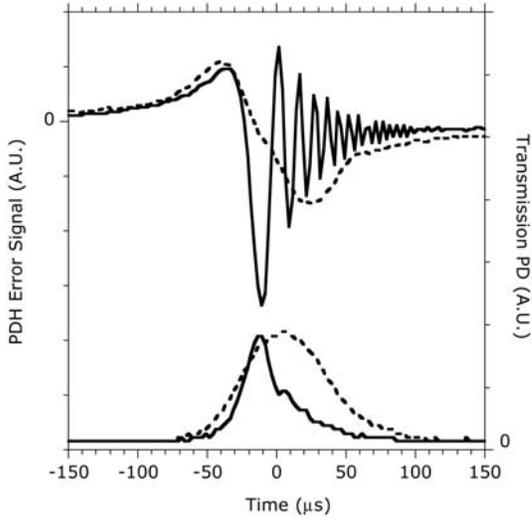
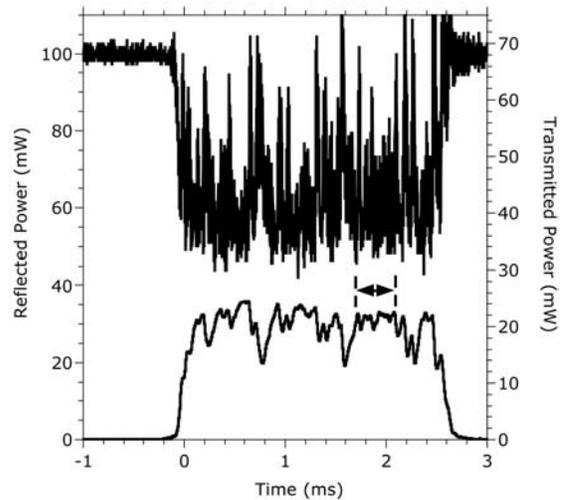


Figure 5.13: Dynamics of cavity response affecting PDH signal. The solid line represents the peak axial mode while the dashed line is measured about 8 axial modes off the peak. At the peak, a transient lock may be susceptible to induced ringing effects which will further destabilize the tracking ability of the servo.

piezo actuator has enough dynamic bandwidth to handle this transient ringing, it is possible that the remaining RMS noise causes net frequency excursions at rates that begin to excite these effects.

Iteratively adjusting the servo eventually enabled the laser to lock in short intervals to the peak axial mode for the narrowest bandwidth cavities tested. Fig. 5.14 displays a high finesse cavity ($\mathcal{F} \approx 12000$) where intermittent locks lasting only a few ms were measured. The achievable peak gains of nearly 5000, however, clearly indicates that cavity enhancement is still unaffected by any aspects of mode-locking.

Figure 5.14: Reflected and transmitted photodiode signals for a laser locked to the peak axial mode in a 6.7 kHz BW cavity. The reflected signal shows large frequency fluctuations although the peak coupling is near predicted values. This cavity was formed from a ST \otimes SF pair where $T_1 = 309$ ppm, $T_2 = 124$ ppm, and the calculated parameters are $b = 1890$ and $c_0 = 0.38$. The added losses are therefore 94 ppm, and the gain ≈ 4500 .



One other interesting test with a narrow-band cavity involves measuring the steady-state transmitted power for a set of sequential axial modes near the peak. The expected values depend only on the cavity bandwidth and the optical pulse length. Fig. 5.15 is a fit of measured transmitted power, normalized to the peak, for a series of axial modes scanned in a single pass. Here, the cavity bandwidth is known, so the fit measures instead the optical pulse length.

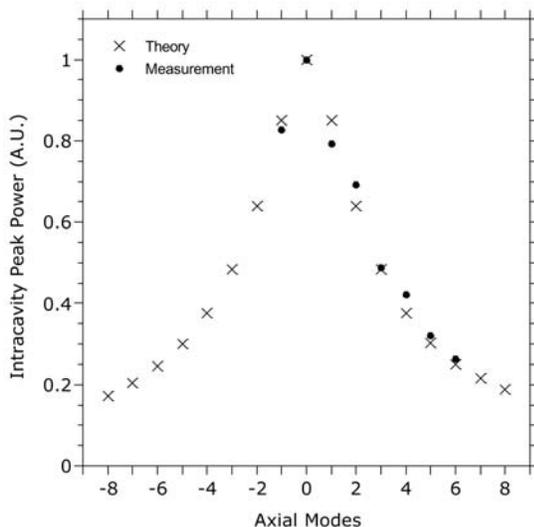


Figure 5.15: Peak stored power measured as a function of axial mode in a 10.4 kHz BW cavity. By properly setting the translation stage velocity on one cavity mirror, the feedback system will re-lock sequentially at each axial mode long enough to measure steady-state transmitted power. A curve fit then measures an optical pulse length of 27 ps FWHM (specified at 25–30 ps).

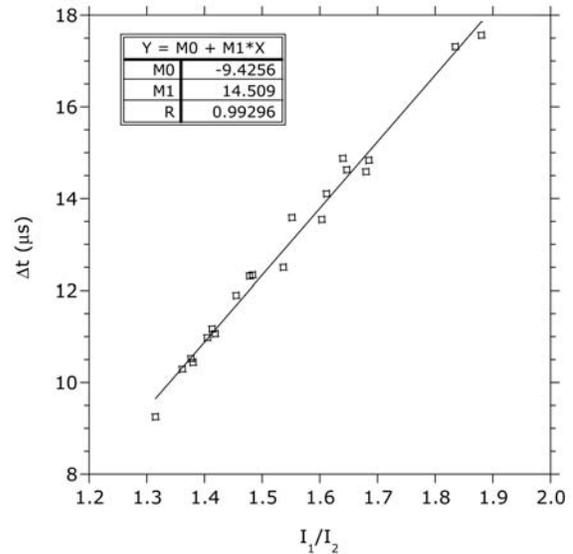
Estimating Cavity Bandwidths Using Transmission Ringing

The dynamics of the transmitted pulse (Sect. 5.2.3) gives another measure of the cavity bandwidth; it has the attractive advantage of not needing two independent and calibrated power measurements, nor any need to lock the laser to the cavity resonance. As a test of the method, the 6.7 kHz bandwidth cavity was scanned at two “fixed” sweep rates and the transmitted pulses recorded. Although the free-running laser added significant uncertainty to the actual frequency sweep rate through resonance, the fitted parameters indeed followed a linear relation. However, fitting the slope of the line predicted a much smaller cavity bandwidth than measured with the usual reflected and transmitted power (Fig. 5.16).

While the predicted losses were too small to be reasonably believed, this measurement still provided a very practical service. When assembling a pair of cavity

mirrors, a relative measure of cavity bandwidth could be quickly ascertained by a few snapshots of the transmission photodiode. If the values were higher than expected, a mirror was replaced and remeasured quickly to find a working set with low losses. This measurement may also be used track the relative performance of a cavity over time, or to monitor power calibration.

Figure 5.16: Fit of transmission ringing coefficients to measure cavity bandwidth. The data represent 20 transmission photodiode pulses where the ratio of the first two intensity peaks is plotted against the relative delay between them. The slope is predicted to be $\tau/2$ (Eq. 5.11). The measured bandwidth is then 5.5 kHz, implying the losses are < 10 ppm (compared to the nominal ≈ 100 ppm).



5.3.4 Discussion

These tests demonstrate that solid-state CW passively mode-locked lasers behave, in practical terms, exactly like single mode CW lasers, or more precisely, a collection of uniformly spaced phase-locked CW lasers in the frequency domain. Hence, all the usual CW cavity power enhancement and stabilization strategies can be successfully applied to these mode-locked lasers with little modification.

The major performance limits, as seen in these experiments, primarily derive from the laser's free-running frequency noise. Compared to the monolithic NPRO cavity, a mode-locked laser cavity is inherently large and more susceptible to coupling acoustic or mechanical noise from the environment to the laser optics. The measured characteristic noise is consistent with this observation, where large frequency excursions at low (kHz) Fourier frequencies dominate.

There are several possible routes to achieve the design goal of a stable cavity power enhancement of 10,000. First, mode-locked lasers are engineered without any consideration of minimizing the free-running frequency noise since very few applications require it; some of this noise will then be passively removed by a more careful mechanical design of the cavity. Secondly, these experiments used the laser's factory installed piezo-ceramic mirror mount, but using a higher bandwidth piezo and lighter mirror would increase the available control bandwidth. To dramatically increase the feedback bandwidth even more, an electro-optic crystal could be used either inside or outside the laser cavity to provide correction for small frequency or phase fluctuations. The larger bandwidth is not necessary to suppress noise at these high frequencies but rather to allow very high gain at the kHz level while producing a suitable phase relationship near the unity gain frequency desired for stable operation. Lastly, one may consider using a staged approach, as done at LIGO [49], in which two or more cavities of decreasing bandwidth sequentially reduce the intensity and frequency noise. Practically, this last option is not as attractive as using a single cavity for both gain enhancement and frequency stabilization.

The ultimate cavity gain enhancement limit will likely have little to do with the laser itself. With appropriately designed servos, the laser noise will be reduced well below the passive cavity bandwidth. As the gain increases, however, the high circulating power will place practical limits on how much optical power density the mirrors can withstand before they are damaged. The power density on the mirrors for the initial X-ray source design parameters is smaller than this expected damage threshold, but, as coating technologies improve, raising the optical circulating power will be a natural upgrade path.

Concluding Remarks

The experimental section of this dissertation addressed a major technical uncertainty of the overall device design. The lack of previous experimental work in pulsed high-gain external cavities is not due to any fundamental technical limitations but simply that no applications had emerged to motivate this study. Combining the work in stabilized CW lasers together with an understanding of mode-locked laser frequency comb stability, external cavities may be successfully designed for high gain optical storage—as well as providing ideal self-references for frequency stabilization techniques. Stored, high-energy, optical pulses may be useful for other physics applications, such as electron beam diagnostics [27], electron beam cooling [26, 42], and pulsed cavity ring-down spectroscopy [17]. This preliminary investigation confirmed the fact that such optical cavities are within technical reach, although more experimental work is needed to find the optimal trade-off of stability vs. gain, especially for a given power-density limit on the cavity mirrors. Maximum power handling, related to materials or coating techniques, also needs further investigation.

In sum, a Thomson scattering device which uses a compact electron storage ring together with an optical pulse gain enhancement cavity is advantageous in design and appears technically feasible. A stored, low emittance electron beam produces nearly monochromatic, collimated X-rays with a source brightness largely inherited from the electron beam emittance. The dynamic stability and performance required of the storage beam is made possible by frequent single-turn, full energy re-injection. At a 100 Hz re-injection rate, the 100 MHz revolution frequency converts an average injector power of a few watts in to a few megawatts of stored beam power. Likewise, the external optical cavity converts a 10 W drive laser into a 100 kW average power

laser in terms of stored optical power. This effective use of electrons and photons is essential to the design of a Thomson scattering source. Furthermore, since the pulses collide head-on, the integrated flux is maximized in this design where every photon has the chance to interact with every electron. No other proposed device in the literature matches both the beam quality and average X-ray intensity as this laser-electron storage ring design. The growing importance of synchrotron radiation together with the predicted performance of this device merits continuing this work.

Bibliography

- [1] A. Abromovici and J. Chapsky. *Feedback Control Systems: A Fast-Track Guide for Scientists and Engineers*. Kluwer Academic Publishers, Mass., USA, 2000.
- [2] D.Z. Anderson. Alignment of resonant optical cavities. *Applied Optics*, 23(17):2944 – 9, September 1984.
- [3] F. Arfelli et al. Mammography with synchrotron radiation: phase-detection techniques. *Radiology*, 215(1):286 – 93, April 2000.
- [4] D. Baker and A. Sali. Protein structure prediction and structural genomics. *Science*, 293(5540):93 – 6, October 2001.
- [5] P. Baumeister. *Optical coating technology : a book used at a five day short course, Engineering 823.17 : presented at University of California at Los Angeles, University Extension, Department of Engineering, Information Systems and Technical Management*. Sholem Press, Sebastopol, CA, 1998.
- [6] F. Bayer-Helms. Coupling coefficients of an incident wave and the modes of a spherical optical resonator in the case of mismatching and misalignment. *Applied Optics*, 23(9):1369 – 80, May 1984.
- [7] J.C. Bienfang, R.F. Teehan, and C.A. Denman. Phase noise transfer in resonant optical cavities. *Review of Scientific Instruments*, 72(8):3208 – 14, August 2001.
- [8] J.D. Bjorken and S. Mtingwa. Intrabeam scattering. *Particle Accelerators*, 13(3-4):115 – 43, 1983.

- [9] E.D. Black. An introduction to Pound-Drever-Hall laser frequency stabilization. *American Journal of Physics*, 69(1):79 – 87, January 2001.
- [10] C.L. Bohn. Coherent synchrotron radiation: theory and experiments. *AIP Conference Proceedings*, (647):81 – 95, 2002.
- [11] E. Bulyak et al. Compact X-ray source based on Compton backscattering. *Nuclear Instruments & Methods in Physics Research, Section A (Accelerators, Spectrometers, Detectors and Associated Equipment)*, 487(3):241 – 8, July 2002.
- [12] P.G. Burkhalter et al. X-ray damage in optical coatings. *Review of Scientific Instruments*, 66(1 pt.2):795 – 7, January 1995.
- [13] Shiuh Chao, Wen-Hsiang Wang, and Cheng-Chung Lee. Low-loss dielectric mirror with ion-beam-sputtered $\text{TiO}_2\text{-SiO}_2$ mixed films. *Applied Optics*, 40(13):2177 – 82, May 2001.
- [14] J. Chen et al. Compact high-brightness radiation sources. *Nuclear Instruments & Methods in Physics Research, Section A (Accelerators, Spectrometers, Detectors and Associated Equipment)*, 358(1-3):14 – 17, April 1995.
- [15] K. Chouffani et al. Laser-Compton scattering from a 20 MeV electron beam. *Nuclear Instruments & Methods in Physics Research, Section A (Accelerators, Spectrometers, Detectors and Associated Equipment)*, 495(2):95 – 106, December 2002.
- [16] D.B. Cline, A. Garren, M. Green, J. Kolonko, and K. Lee. An ultra-compact hard X-ray superconducting light source for biotechnology and industrial use. *Nuclear Instruments & Methods in Physics Research, Section B (Beam Interactions with Materials and Atoms)*, 139(1-4):531 – 6, April 1998.
- [17] E.R. Crosson et al. Pulse-stacked cavity ring-down spectroscopy. *Review of Scientific Instruments*, 70(1 pt.1):– 2, 4–10, January 1999.

- [18] A. Cutolo and S. Solimeno. Alignment and performance of almost concentric resonators for low gain free-electron lasers. *Applied Optics*, 26(1):52 – 62, January 1987.
- [19] T. Day. *Frequency stabilized solid state lasers for coherent optical communications*. PhD thesis, Stanford University, Stanford, CA, Sept 1990.
- [20] A.M. De Riva et al. Very high Q frequency-locked Fabry-Perot cavity. *Review of Scientific Instruments*, 67(8):2680 – 4, August 1996.
- [21] S.A. Diddams et al. Direct RF to optical frequency measurements with a femtosecond laser comb. *IEEE Transactions on Instrumentation and Measurement*, 50(2):552 – 5, April 2001.
- [22] S. Doniach. Fourth-generation X-ray sources: some possible applications to biology. *Journal of Synchrotron Radiation*, 7(pt.3):116 – 20, May 2000.
- [23] J. Drenth. *Principles of Protein X-Ray Crystallography. Second edition*. Springer-Verlag, New York, USA, 1999.
- [24] R.W.P. Drever et al. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B (Photophysics and Laser Chemistry)*, B31(2):97 – 105, June 1983.
- [25] R.C. Elton et al. Time-resolved measurements of X-ray damage to optical coatings. *Journal of Applied Physics*, 81(3):1184 – 91, February 1997.
- [26] E. Esarey. Laser cooling of electron beams via Thomson scattering. *Nuclear Instruments & Methods in Physics Research, Section A (Accelerators, Spectrometers, Detectors and Associated Equipment)*, 455(1):7 – 14, November 2000.
- [27] N. Falletto et al. Compton scattering off polarized electrons with a high-finesse Fabry-Perot cavity at JLab. *Nuclear Instruments & Methods in Physics Research, Section A (Accelerators, Spectrometers, Detectors and Associated Equipment)*, 459(3):412 – 25, March 2001.

- [28] T.M. Fortier, D.J. Jones, Jun Ye, S.T. Cundiff, and R.S. Windeler. Long-term carrier-envelope phase coherence. *Optics Letters*, 27(16):1436 – 8, August 2002.
- [29] G. Franklin, J. Powell, and A. Emami-Naeini. *Feedback control of dynamic systems*. Prentice Hall, New Jersey, 2002.
- [30] M.A. Furman. Hourglass effects for asymmetric colliders. In *IEEE Particle Accelerator Conference. Accelerator Science and Technology*, San Francisco, CA, USA, May 1991.
- [31] M.A. Furman and M.S. Zisman. Luminosity. In A. W. Chao and M. Tigner, editors, *Handbook of Accelerator Physics and Engineering.*, chapter 4.1, pages 247 – 251. World Scientific, 1999.
- [32] E. Garman and C. Nave. Radiation damage to crystalline biological molecules: current view. *Journal of Synchrotron Radiation*, 9(6):327 – 328, November 2002.
- [33] D. Garzella et al. Mirror degradation and heating in storage ring FELs. *Nuclear Instruments & Methods in Physics Research, Section A (Accelerators, Spectrometers, Detectors and Associated Equipment)*, 358(1-3):387 – 91, April 1995.
- [34] A. Gatto et al. Toward resistant UV mirrors at 200 nm for free electron lasers: manufacture characterizations and degradation tests. *Proceedings of the SPIE - The International Society for Optical Engineering*, 4102:261 – 75, 2000.
- [35] A. Gatto et al. High-performance deep-ultraviolet optics for free-electron lasers. *Applied Optics*, 41(16):3236 – 41, June 2002.
- [36] B.I. Grishanov, F.V. Podgorny, J. Ruemmler, and V.D. Shiltsev. Very fast kicker with high repetition rate for accelerator applications. *Nuclear Instruments & Methods in Physics Research, Section A (Accelerators, Spectrometers, Detectors and Associated Equipment)*, 396(1-2):28 – 34, September 1997.
- [37] P. D. Haar. *Pulse stacking in the Stanford external cavity and photo-induced reflectivity in the infrared*. PhD thesis, Stanford University, Stanford, CA, Dec 1996.

- [38] R. Hauck, H.P. Kortz, and H. Weber. Misalignment sensitivity of optical resonators. *Applied Optics*, 19(4):598 – 601, February 1980.
- [39] F.W. Helbing, G. Steinmeyer, J. Stenger, H.R. Telle, and U. Keller. Carrier-envelope-offset dynamics and stabilization of femtosecond pulses. *Applied Physics B (Lasers and Optics)*, B74(suppl.):S35 – 42, June 2002.
- [40] A. Hofmann. Theory of synchrotron radiation. Technical Report ACD-Note 38, SSRL, Menlo Park, CA, Sept 1986.
- [41] I.C. Hsu, Cha-Ching Chu, and Chuan-Ing Yu. Energy measurement of relativistic electron beams by laser Compton scattering. *Physical Review E (Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics)*, 54(5):5657 – 63, November 1996.
- [42] Z. Huang. *Radiative cooling of relativistic electron beams*. PhD thesis, Stanford University, Stanford, CA, May 1998.
- [43] Z. Huang and R.D. Ruth. Laser-electron storage ring. *Physical Review Letters*, 80(5):976 – 9, February 1998.
- [44] Z. Huang and R.D. Ruth. Radiative cooling of relativistic electron beams. In *1999 Particle Accelerator Conference*, New York, NY, USA, March 1999.
- [45] Y. Hwu, Wen-Li Tsai, A. Groso, G. Margaritondo, and Jung Ho Je. Coherence-enhanced synchrotron radiology: simple theory and practical applications. *Journal of Physics D (Applied Physics)*, 35(13):R105 – 20, July 2002.
- [46] J. D. Jackson. *Classical electrodynamics. Third edition*. Wiley, New York, 1999.
- [47] C. Janke. Thermal loading of optical components in interferometric systems. Master’s thesis, Fachhochschule O/O/W and Stanford University, Emden, Germany and Stanford, CA, Dec 2000.
- [48] T.F. Johnston, Jr. Beam propagation (M^2) measurement made as easy as it gets: The four-cuts method. *Applied Optics*, 37(21):4840 – 50, July 1998.

- [49] S. Kawamura, A. Abramovici, and M.E. Zucker. Improved multistage wide band laser frequency stabilization [for LIGO interferometer]. *Review of Scientific Instruments*, 68(1 pt.1):223 – 9, January 1997.
- [50] V.C. Kempson et al. The HELIOS series of advanced synchrotron based X-ray sources. *Journal de Physique III (Applied Physics, Materials Science, Fluids, Plasma and Instrumentation)*, 4(9):1679 – 86, September 1994.
- [51] M.Kh. Khokonov and Jr. Carrigan, R.A. The relationship of channeling radiation to Thomson scattering and the relative efficiency of X-ray production by intense electron beams. *Nuclear Instruments & Methods in Physics Research, Section B (Beam Interactions with Materials and Atoms)*, 145(1-2):133 – 41, October 1998.
- [52] H. Kotaki et al. Compact X-ray sources by intense laser interactions with beams and plasmas. *Nuclear Instruments & Methods in Physics Research, Section A (Accelerators, Spectrometers, Detectors and Associated Equipment)*, 455(1):166 – 71, November 2000.
- [53] M.J. Lawrence, B. Willke, M.E. Husman, E.K. Gustafson, and R.L. Byer. Dynamic response of a Fabry-Perot interferometer. *Journal of the Optical Society of America B (Optical Physics)*, 16(4):523 – 32, April 1999.
- [54] R. Li. The impact of coherent synchrotron radiation on the beam transport of short bunches. In *1999 Particle Accelerator Conference*, New York, NY, USA, March 1999.
- [55] Y. Li, Z. Huang, M.D. Borland, and S. Milton. Small-angle Thomson scattering of ultrafast laser pulses for bright sub-100-fs X-ray radiation. *Physical Review Special Topics-Accelerators and Beams*, 5(4), April 2002.
- [56] J.P. Lidestri. A synchrotron design for macromolecular structure determination. In *2001 Particle Accelerator Conference (PAC2001)*, Chicago, IL, USA, June 2001.

- [57] J. Macken and R. Salvage. Alignment-insensitive resonators using focusing corner cubes. *Proceedings of the SPIE - The International Society for Optical Engineering*, 1224:386 – 92, 1990.
- [58] H.A. Macleod. *Thin-film optical filters. Third edition.* IOP Publishing, Bristol, UK, 2001.
- [59] G. McConnell, A.I. Ferguson, and N. Langford. Cavity-augmented frequency tripling of a continuous wave mode-locked laser. *Journal of Physics D (Applied Physics)*, 34(16):2408 – 13, August 2001.
- [60] O. Mor and A. Arie. Performance analysis of Drever-Hall laser frequency stabilization using a proportional+integral servo. *IEEE Journal of Quantum Electronics*, 33(4):532 – 40, April 1997.
- [61] K. Nakajima et al. A short pulse X-ray generation via Thomson scattering of ultrashort laser pulses by relativistic electron beams. In *1997 Particle Accelerator Conference*, Vancouver, BC, Canada, May 1997.
- [62] C. Nave. Source optical and detector requirements for X-ray diffraction and scattering. *Journal of Synchrotron Radiation*, 5(pt.3):645 – 7, May 1998.
- [63] C. Nave. Matching X-ray source, optics and detectors to protein crystallography requirements. *Acta Crystallographica Section D*, 55(10):1663 – 1668, October 1999.
- [64] M.A. Persaud, J.M. Tolchard, and A.I. Ferguson. Efficient generation of picosecond pulses at 243 nm. *IEEE Journal of Quantum Electronics*, 26(7):1253 – 8, July 1990.
- [65] A. Piwinski. Touscheck effect and intrabeam scattering. In A. W. Chao and M. Tigner, editors, *Handbook of Accelerator Physics and Engineering.*, chapter 2.5.9, pages 125 – 127. World Scientific, 1999.

- [66] J. Poirson, F. Bretenaker, M. Vallet, and A. Le Floch. Analytical and experimental study of ringing effects in a Fabry-Perot cavity. Application to the measurement of high finesses. *Journal of the Optical Society of America B (Optical Physics)*, 14(11):2811 – 17, November 1997.
- [67] Ch. Salomon, D. Hils, and J.L. Hall. Laser stabilization at the millihertz level. *Journal of the Optical Society of America B (Optical Physics)*, 5(8):1576 – 87, August 1988.
- [68] E.S. Sarachik and G.T. Schappert. Classical theory of the scattering of intense laser radiation by free electrons. *Physical Review D (Particles and Fields)*, 1(10):2738 – 53, May 1970.
- [69] A. Schoof, J. Grunert, S. Ritter, and A. Hemmerich. Reducing the linewidth of a diode laser below 30 Hz by stabilization to a reference cavity with a finesse above 10^5 . *Optics Letters*, 26(20):1562 – 4, October 2001.
- [70] R.E. Sherriff. Analytic expressions for group-delay dispersion and cubic dispersion in arbitrary prism sequences. *Journal of the Optical Society of America B (Optical Physics)*, 15(3):1224 – 30, March 1998.
- [71] A. E. Siegman. *Lasers*. University Science Books, Mill Valley, Calif., 1986.
- [72] T.B. Simpson, T. Day, F. Doft, M.M. Malley, and G.W. Sutton. Frequency-stabilized mode-locked solid-state laser system for precision range-Doppler imaging. *IEEE Journal of Quantum Electronics*, 29(9):2489 – 96, September 1993.
- [73] D.A. Smith and D.I. Shernoff. Simple measurement of gain and loss in ultralow loss optical resonators. *Applied Optics*, 24(12):1722 – 3, June 1985.
- [74] T.I. Smith, P. Haar, and H.A. Schwettman. Pulse stacking in the SCA/FEL external cavity. *Nuclear Instruments & Methods in Physics Research, Section A (Accelerators, Spectrometers, Detectors and Associated Equipment)*, 393(1-3):245 – 51, July 1997.

- [75] P. Sprangle, A. Ting, E. Esarey, and A. Fisher. Tunable short pulse hard X-rays from a compact laser synchrotron source. *Journal of Applied Physics*, 72(11):5032 – 8, December 1992.
- [76] K.G. Steffen. *High-Energy Beam Optics*. Interscience Publishers, John Wiley & Sons, New York, 1965.
- [77] V. Telnov. Laser cooling of electron beams at linear colliders. *Nuclear Instruments & Methods in Physics Research, Section A (Accelerators, Spectrometers, Detectors and Associated Equipment)*, 455(1):80 – 9, November 2000.
- [78] Y.V. Troitski. Synthesis of mirrors with anomalous dispersion of the reflection phase. *Optical Engineering*, 34(5):1503 – 7, May 1995.
- [79] A. Tsunemi et al. Ultra-bright X-ray generation using inverse Compton scattering of picosecond CO₂ laser pulses. In *1999 Particle Accelerator Conference*, New York, NY, USA, March 1999.
- [80] T. Udem, J. Reichert, R. Holzwarth, and T.W. Hansch. Accurate measurement of large optical frequency differences with a mode-locked laser. *Optics Letters*, 24(13):881 – 3, July 1999.
- [81] A. Ueda et al. Ultra-high quality cavity with 1.5 ppm loss at 1064 nm. *Optical Review*, 3(5):369 – 72, September 1996.
- [82] N. Uehara et al. Ultralow-loss mirror of the parts-in-10⁶ level at 1064 nm. *Optics Letters*, 20(6):530 – 2, March 1995.
- [83] N. Uehara and K. Ueda. 193-mhz beat linewidth of frequency-stabilized laser-diode-pumped nd:yag ring lasers. *Optics Letters*, 18(7):505 – 7, April 1993.
- [84] N. Uehara and K. Ueda. Accurate measurement of the radius of curvature of a concave mirror and the power dependence in a high-finesse Fabry-Perot interferometer. *Applied Optics*, 34(25):5611 – 19, September 1995.

- [85] J. Urakawa et al. Laser super cavity. In *21st ICFA Beam Dynamics Workshop on Laser-Beam Interactions*, Stony Brook, NY, USA, June 2001.
- [86] M. Venturini. Scaling of third-order quadrupole aberrations with fringe field extension. In *1999 Particle Accelerator Conference, 27 March-2 April 1999, New York, NY, USA*.
- [87] M. Venturini. Study of intrabeam scattering in low-energy electron rings. In *IEEE Particle Accelerator Conference (PAC2001)*, Chicago, IL, USA, Jun 2001.
- [88] F. Verluise, V. Laude, J.-P. Huignard, P. Tournois, and A. Migus. Arbitrary dispersion control of ultrashort optical pulses with acoustic waves. *Journal of the Optical Society of America B (Optical Physics)*, 17(1):138 – 45, January 2000.
- [89] H. Wiedemann. *Particle Accelerator Physics I : Basic Principles and Linear Beam Dynamics*. Springer Verlag, New York, 1998.
- [90] H. Wiedemann. *Particle Accelerator Physics II 2. ed.. Nonlinear and Higher-Order Beam Dynamics*. Springer, Berlin, Germany, 1999.
- [91] H. Winick. Synchrotron radiation sources: Past, present and future. In A. Haase, G. Landwehr, and E. Umbach, editors, *Röntgen Centennial. X-rays in Natural and Life Sciences.*, pages 45 – 64. World Scientific, 1997.
- [92] Yang Xiang. Focus retrocollimated interferometry for long-radius-of-curvature measurement. *Applied Optics*, 40(34):6210 – 14, December 2001.
- [93] H. Yamada et al. Development of the hard X-ray source based on a table-top electron storage ring. *Nuclear Instruments & Methods in Physics Research, Section A (Accelerators, Spectrometers, Detectors and Associated Equipment)*, 467-468(pt.1):122 – 5, July 2001.
- [94] K. Yamada et al. Degradation and restoration of dielectric-coated cavity mirrors in the NIJI-IV FEL. *Nuclear Instruments & Methods in Physics Research, Section A (Accelerators, Spectrometers, Detectors and Associated Equipment)*, 358(1-3):392 – 5, April 1995.

- [95] M. Yasumoto, T. Tomimasu, S. Nishihara, and N. Umesaki. Optical damage of multi-layer mirrors for UV-FEL resonator induced with intense pico-second pulse FELs and Nd:YLF lasers. *Optical Materials*, 15(1):59 – 63, September 2000.