Chapter

5.3

Page 1:

Welcome, Dear Students of Dr M A Gondal, everyone. Today, we embark on a detailed exploration of Chapter 5, Section 3, focusing on the "Spectral Characteristics of Laser Emission." This is a cornerstone topic in understanding how lasers actually produce the light we observe and utilize. We'll be delving into the intricacies that determine the precise colors, or more accurately, the frequencies, that a laser emits. The material we'll cover has been prepared by Distinguished Professor Doctor M A Gondal for our Physics 608 Laser Spectroscopy course here at KFUPM. Let's dive in.

Page 2:

Alright, let's set the stage for our discussion on the "Spectral Characteristics of Laser Emission." The primary goal for this block of slides, and indeed this section of the course, is to build a rigorous, step-by-step understanding of a fundamental question: How do an active medium and an optical resonator, working in concert, determine the *exact set of frequencies* emitted by a laser? It's not enough to say a laser emits red light; we need to understand *which* red frequencies, how narrow they are, and what physical processes dictate this specificity. This understanding is absolutely critical for anyone designing, using, or interpreting experiments involving lasers, especially in high-resolution spectroscopy.

To achieve this goal, we need to dissect the problem into its key ingredients. These are components that we must analyze separately first, to understand their individual contributions, and then combine them to see the full picture. This approach of breaking down a complex system is a hallmark of physics.

The first crucial ingredient is "The spectral gain profile of the amplifying transition in the medium." Let's unpack this. The "active medium," as you know, is where light amplification occurs. This amplification is tied to a specific "amplifying transition" – an electronic, vibrational, or rotational transition within the atoms or molecules of the medium that has a population inversion. This transition doesn't provide equal gain at all frequencies; instead, it has a "spectral gain profile." This profile is a curve, plotting gain as a function of frequency. It has a certain central frequency and a certain width, determined by the physics of the gain medium itself, such as Doppler broadening or lifetime broadening. Understanding the shape and characteristics of this gain profile is our first step because the laser can only oscillate at frequencies where there is sufficient gain to overcome losses.

Page 3:

Continuing with the key ingredients that determine laser emission frequencies:

The second ingredient is "The discrete resonance spectrum (axial and transverse modes) of the passive optical cavity." A "passive optical cavity" typically consists of two mirrors facing each other. It's called "passive" because, at this stage of analysis, we're not yet considering the gain medium inside it. Such a cavity doesn't support any arbitrary frequency of light. Instead, it only allows standing waves to form at specific, discrete resonant frequencies. These are the "eigenfrequencies" of the cavity. These resonant frequencies are determined by the cavity's geometry, primarily its length. We'll see that there are "axial modes," which relate to standing waves along the main axis between the mirrors, and "transverse modes" (often denoted TEM$_{mn}$ modes), which describe the field distribution in the plane perpendicular to the axis. The interplay between this discrete set of allowed cavity frequencies and the continuous gain profile of the active medium is central to laser operation.

The third ingredient involves "The non-linear, intensity-dependent phenomena that arise once oscillation starts." This is where things get really interesting and often more complex. Once the laser begins to oscillate, the intensity of the light inside the cavity can become very high.

This high intensity can, in turn, modify the properties of the active medium and the interaction. We're talking about phenomena such as:

"Gain saturation": As the intensity builds up, it depletes the population inversion, thereby reducing the gain. This is a crucial self-regulating mechanism in lasers.

"Mode competition": If multiple cavity modes fall under the gain profile, they might compete for the available gain. Sometimes one mode wins out, leading to single-mode operation; other times, multiple modes can coexist.

"Spatial hole burning": In standing-wave cavities, the intensity pattern has nodes and antinodes. Saturation can occur preferentially at the antinodes, "burning holes" in the spatial distribution of the gain. This can allow other modes, whose antinodes are in different locations, to lase.

"Mode pulling": The refractive index of the gain medium itself can be frequency-dependent, especially near the gain line center. This can slightly shift, or "pull," the actual lasing frequencies away from the passive cavity resonance frequencies.

These non-linear effects are critical for understanding the stable operating characteristics of a real laser.

So, what's our "Strategy for the forthcoming slides"? We'll tackle this systematically:

1. First, we will "Introduce mathematical description of the passive cavity." We need to understand the allowed frequencies of an empty resonator. This provides the basic framework, the "comb" of possible frequencies.

2. Second, we will "Insert an active medium, derive modified eigenfrequencies." We'll see how placing an active medium, with its own refractive index, inside the cavity changes the optical path length and thus shifts these eigenfrequencies.

Page 4:

Continuing with our strategy:

3. Third, we will "Develop the laser threshold condition from first principles." This is a pivotal concept. The threshold condition defines the minimum gain required from the active medium to overcome all the losses in the cavity (like mirror transmission, absorption, scattering) for laser oscillation to begin. It tells us *when* a laser will lase.

4. And fourth, we will "Examine line-narrowing effects, saturation physics, multimode behaviour, and the frequency shifts produced by dispersion in the gain medium." This involves looking at how the interaction of the gain profile and the cavity modes, along with saturation, leads to the actual observed laser spectrum. Line-narrowing is a key characteristic of laser light, and we'll explore why laser output is much narrower than the gain profile or the cavity resonances alone. We'll delve deeper into saturation physics, how multiple modes behave and interact, and quantitatively assess the frequency shifts due to dispersion – that's the mode pulling we mentioned.

Now, to ensure clarity and consistency, let's establish some "Conventions adopted throughout" this discussion:

* The symbol c (lowercase cee) will represent the speed of light in vacuum. Its units will typically be meters per second, written as m s^{-1}. The value is approximately 3×10^8 meters per second.

* The symbol d (lowercase dee) will represent the physical mirror separation in a linear cavity. Its units will be meters, denoted as m.

Page 5:

Continuing with our conventions:

* The symbol $n(\nu)$ – that's lowercase n as a function of ν (the Greek letter ν, representing frequency) – will denote the frequency-dependent refractive index of the active medium. The refractive index, as you know, is a

dimensionless quantity. The fact that it's frequency-dependent, $n(\nu)$, especially near the atomic or molecular transition providing gain, is crucial for understanding phenomena like mode pulling.

* The symbol L (capital Ell) will represent the length of the gain medium inside the cavity. This is distinct from d, the total mirror separation. L will also be in units of meters, m (em).

* The "Axial mode index q" (lowercase cue) will be an integer, belonging to the set of integers denoted by the symbol $\mathbb{Z}$ (from the German 'Zahlen'). So, q can be zero, one, two, three, and so on. It quantifies how many half-wavelengths of the light fit into the cavity for a particular axial mode.

A quick note on notation: "All symbols will be re-defined on the..."

Page 6:

"...Slide where they first appear." This is just a pedagogical commitment to ensure clarity. Even if we introduce a symbol now, when it becomes central to an equation or a discussion on a later slide, we'll briefly remind you of its meaning.

Finally, regarding units: "CGS and SI units will be stated explicitly when needed." Physics relies on consistent use of units, and while SI (Système International) units are generally preferred, sometimes historical context or convenience in certain subfields might lead to CGS (Centimeter-Gram-Second) units. We will be explicit to avoid any confusion.

Now that we've set the stage and defined our initial conventions, let's move to our first main topic: the passive resonator.

Page 7:

Alright, let's begin our quantitative journey with a "Passive Fabry-Perot Resonator — Eigenfrequency Review." The Fabry-Perot resonator, or etalon, typically consisting of two parallel mirrors, is the most fundamental type of optical resonator. We call it "passive" here because we are considering it without any gain medium inside – just the mirrors and the space between them. This review will establish the baseline for the frequencies that *could* resonate if there were light present.

* First, "Consider two perfectly parallel mirrors separated by distance d." Imagine these two mirrors facing each other. The "perfectly parallel" is an idealization, but it simplifies the initial analysis. d is the physical separation. "For vacuum between the mirrors, the round-trip optical path length is $2\,d$." This is straightforward: light travels a distance d from one mirror to the other, reflects, and travels back a distance d, completing a round trip of $2\,d$ (two dee).

* Now, for light to resonate within this cavity, it must form a standing wave. The "Standing-wave condition" requires that the electric field must be zero at the surface of perfectly conducting mirrors. This leads to the condition that the cavity length d must be an integer multiple of half-wavelengths. Equivalently, for a round trip, the phase change must be an integer multiple of 2π. This leads to the equation shown:

$$2\,d = q \times \frac{c}{\nu_{q,\text{vac}}}$$

Let's break this down: "Two dee equals cue times the ratio of cee over nu sub cue comma vac." Here, $2\,d$ (two dee) is the round-trip path length in meters. q (cue) is the axial mode index, an integer $(0,1,2,\dots)$, as we defined earlier. It essentially counts the number of half-wavelengths that fit into the cavity length d. c (cee) is the speed of light in vacuum, in meters per second. $\nu_{q,\text{vac}}$ (nu sub cue comma vac) is the eigenfrequency of the q-th axial mode in vacuum, in Hertz. This is the specific frequency that satisfies the standing wave condition for that particular integer q. This equation fundamentally states that for a standing wave to exist, the round-trip path length ($2\,d$) must be an integer multiple (q) of the wavelength ($\lambda = \frac{c}{\nu_{q,\text{vac}}}$).

This condition quantizes the frequencies that can exist stably within the passive cavity.

Page 8:

Continuing from the standing-wave condition, let's define the terms more formally:

$\nu_{q,\text{vac}}$ (nu sub q comma vac): This is the "eigenfrequency of the q-th axial mode in vacuum," and its units are Hertz (Hz). These are the specific, discrete frequencies that "fit" perfectly into the cavity, forming standing waves.

q (lowercase q): This is the "integer mode number," taking values like 0, 1, 2, and so on. Each value of 'q' corresponds to a different axial mode, a different number of half-wavelengths across the cavity length.

Now, we can "Solve for the eigenfrequencies" by simply rearranging the previous equation. If we make $\nu_{q,\text{vac}}$ the subject, we get:

$$\nu_{q,\text{vac}} = \frac{q \cdot c}{2 \cdot d}$$

Let's verbalize this: "nu sub q comma vac equals q times c, all divided by two times d."

This equation is fundamental. It tells us that the allowed frequencies in a passive vacuum cavity are directly proportional to the mode number 'q' and the speed of light c, and inversely proportional to twice the cavity length d. This means the resonant frequencies form an equally spaced ladder or comb of frequencies.

From this, we can define an "Important derived quantity – Free Spectral Range (FSR)." The FSR is a characteristic parameter of any Fabry-Perot cavity.

Page 9

The Free Spectral Range, or FSR, is defined as the spacing between adjacent axial modes. Let's denote it by $\delta\nu$ (delta nu). It's the difference between the frequency of mode $q+1$ and mode q. So,

$$\delta\nu = \nu_{q+1,\text{vac}} - \nu_{q,\text{vac}}$$

Using our formula for $\nu_{q,\text{vac}}$:

$$\delta\nu = \frac{(q+1)c}{2\,d} - \frac{q\,c}{2\,d}$$

Factoring out $\frac{c}{2\,d}$, we get:

$$\delta\nu = \left(\frac{c}{2\,d}\right)(q+1-q)$$

Which simplifies to:

$$\delta\nu = \frac{c}{2\,d}$$

* So, "FSR is the constant spacing between adjacent axial modes of the passive cavity." This is a very important result. It means if you know the cavity length 'd', you immediately know the frequency separation between its resonant modes in vacuum. For example, a 1-meter long cavity ($d = 1\,\text{m}$) would have a $2\,d$ of 2 meters. If c is $3 \times 10^8\,\text{m/s}$, then FSR = $\frac{3\times10^8\,\text{m/s}}{2\,\text{m}} = 1.5 \times 10^8\,\text{Hz}$, or 150 Megahertz. This is a typical FSR for a meter-scale cavity.

* The next point here is about the "Passive cavity half-width (power) of each resonance." Each of these resonant frequencies is not infinitely sharp. They have a finite linewidth, or half-width, often denoted $\Delta\nu_{\text{c}}$ (Delta nu sub c) or related to the cavity finesse. This width is "caused by mirror transmission T (capital Tee) and additional losses γ (gamma)." These losses represent energy escaping the cavity per round trip. The derivation of this half-width, and the related concept of cavity Q-factor and finesse, "will be derived later when $G(\nu)$ is introduced." $G(\nu)$ (capital Gee of nu) will

be our round-trip gain factor, and understanding losses is crucial for that. For now, just appreciate that these resonances have a finite width.

Page 10:

Now we move to a slightly more complex scenario, titled "Inserting an Active Medium — Effective Optical Length." Up until now, we've considered a vacuum between the mirrors. What happens when we place our gain medium, which is a dielectric material, inside this cavity?

* The first bullet point states: "Whenever a dielectric slab of physical length L and frequency-dependent index $n(\nu)$ is placed between the mirrors, the effective round-trip optical path length changes." This is the key insight. Light travels slower in a medium with refractive index greater than one. The optical path length is not just the physical length, but the physical length multiplied by the refractive index. Since n can be frequency-dependent, the optical path length also becomes frequency-dependent.

Therefore, we need to "Replace the earlier geometric length d by an effective length \(d^(\nu)\)." This $d^*(\nu)$ will account for the presence of the dielectric.

Let's look at the "Detailed derivation" of this effective optical length for one pass through the cavity (the slide implies round-trip, but the \(d^*(\nu)\) is usually defined for one pass, and then the round trip is \(2\,d*^(\nu)\). Let's clarify as we go, the equations will make it clear. The later equation \(2\,d^*(\nu)\) confirms \(d*^(\nu)\) is one-way effective length).

1. "Segment outside the gain medium:" If the total mirror separation is d, and the gain medium has a physical length L, then the remaining length inside the cavity is $d - L$. Assuming this part is vacuum (or air, with $n \approx 1$), its optical length is simply its physical length.

Page 11:

So, following the derivation:

The physical length $d - L$ in vacuum contributes an optical length of $d - L$. This is step 1.

2. "Segment inside the gain medium:"

The gain medium itself has a physical length L. Because it has a refractive index $n(\nu)$, this physical length L "gives optical length $n(\nu)L$". This is because the light effectively travels n times further, or takes n times longer, to traverse this physical distance compared to vacuum.

3. "Total optical length:"

So, for a single pass from one mirror to the other, the total effective optical length, $d^*(\nu)$, is the sum of these two parts:

$$d^*(\nu) = (d - L) + n(\nu)L$$

This is the effective one-way optical path length from one mirror to the other when a medium of length L and refractive index $n(\nu)$ is present.

4. "Rearranged form:"

We can rewrite this by factoring out L from the last two terms, or rather by regrouping:

$$d^*(\nu) = d - L + n(\nu)L = d + (n(\nu)L - L)$$

So, $d^*(\nu) = d + [n(\nu) - 1]L$.

This form is quite intuitive. It says the effective length is the original physical length d, plus an extra term $(n(\nu) - 1)L$. This extra term represents the additional optical path introduced by the medium compared to if that same length L were vacuum. If $n(\nu)$ is greater than 1, this term is positive, and the effective optical length is longer than d.

With this new effective optical length \(d^(\nu)\), "The new eigenfrequency condition becomes:" We simply replace d in our previous vacuum eigenfrequency condition with $d^*(\nu)$.

Page 12:

The new eigenfrequency condition, analogous to $2\,d = q \cdot \frac{c}{\nu_{q,\text{vac}}}$ for the vacuum case, now becomes:

$$2 \cdot d^*(\nu) = q \cdot \frac{c}{\nu_{q,\text{act}}}$$

Here, \(d^*(\nu)\) is the frequency-dependent effective* one-way *optical path length we just derived. So, \(2 \cdot d*^(\nu)\) is the effective *round-trip* optical path length. q is still our integer axial mode number. c is the speed of light in vacuum. And critically, $\nu_{q,\text{act}}$ denotes the active-cavity eigenfrequency. These are the resonant frequencies when the active medium is present.

Notice a subtlety here: d^* depends on frequency ν, and the resonant frequency $\nu_{q,\text{act}}$ itself is what we are solving for. This means the equation

$$\frac{2\,d^*\left(\nu_{q,\text{act}}\right)\nu_{q,\text{act}}}{c} = q$$

is now an implicit equation for $\nu_{q,\text{act}}$ because d^* contains $n\left(\nu_{q,\text{act}}\right)$. This is more complex than the vacuum case where d was a constant. This frequency dependence of $n(\nu)$ is the origin of dispersion effects like mode pulling, which we'll discuss in detail later. For now, this equation defines the new set of resonant frequencies.

Page 13:

This page provides a very helpful visual for what we've just discussed, under the title "Chapter 5.3: Spectral Characteristics of Laser Emission" and "Figure: Cavity Configurations and Optical Paths." We have two panels.

Panel (a) is titled "Empty Cavity." It depicts two thick vertical black lines representing the mirrors. The distance between them is labeled d with a double-headed arrow. A dashed red line, representing a light ray, traverses the cavity from left to right, labeled "Optical path = d". This illustrates the simple case: for an empty cavity (vacuum or air with $n \approx 1$), the one-way optical path length is just the physical separation d.

Panel (b) is titled "Cavity with Active Medium." Again, we see the two mirrors, also separated by the physical distance d. However, now, a significant portion of the space between the mirrors is occupied by a light blue rectangle, which is labeled "Active Medium." This active medium has a physical length L, indicated by a double-headed arrow below it. Inside this active medium, the dashed red line representing the light ray's path is labeled "Optical path = $n(\nu)L$" (en of nu times Ell), where $n(\nu)$ is also shown within the blue block. The part of the cavity *not* filled by the active medium is indicated below the entire setup as "Total optical path in vacuum = $d - L$" (dee minus Ell), which refers to the sum of vacuum path segments if the medium isn't filling the whole cavity. Here, it seems the medium of length L is placed somewhere within d.

The crucial formula we derived,

$$d^*(\nu) = (d - L) + n(\nu)L$$

assumes the remaining $d - L$ is vacuum. This diagram beautifully illustrates that the total optical path for one pass is no longer just d, but is modified by the presence of the medium with its refractive index $n(\nu)$ over the length L. The term "Total optical path in vacuum = $d - L$" on the diagram specifically refers to the portion of the cavity *not* occupied by the medium. So, the total one-way optical path is indeed $(d - L) + n(\nu)L$.

This figure clearly contrasts the two scenarios and helps to visualize why the effective optical length d^* must be used when an active medium is present.

Page 14:

Now we encounter a very important phenomenon: "Anomalous Dispersion and Its Spectroscopic Importance."

The first point states: "Within the gain bandwidth of an inverted transition, the real part of the refractive index exhibits anomalous dispersion." Let's break this down. "Gain bandwidth" is the range of frequencies over which the active medium can provide amplification. An "inverted transition" refers to a pair of energy levels where the upper level is more populated than the lower, a necessary condition for gain. "The real part of the refractive index," $n(\nu)$, is what affects the phase velocity of light and thus the optical path length. "Anomalous dispersion" is a specific behavior of the refractive index with frequency. Normally, for transparent materials far from resonance, refractive index increases with frequency (normal dispersion, like in a prism). Anomalous dispersion refers to regions where the derivative $\frac{dn}{d\nu}$ is negative, meaning the refractive index *decreases* with increasing frequency. This typically occurs in the vicinity of an absorption line. For an *inverted* transition (a gain line), the behavior is related but, as we'll see, leads to characteristic shifts.

The slide specifies the behavior of $n(\nu)$ near the center of the gain line, let's call it ν_0:

" $n(\nu)$ decreases below unity on the low-frequency side ($\nu < \nu_0$)." That is, for frequencies slightly less than the line center frequency ν_0, the refractive index can actually be less than 1.

" $n(\nu)$ increases above unity on the high-frequency side ($\nu > \nu_0$)." For frequencies slightly greater than ν_0, the refractive index is greater than 1.

This "S"-shaped curve for $n(\nu)$ around ν_0, particularly the region of rapid change, is characteristic of anomalous dispersion associated with a gain feature. The fact that $n(\nu)$ can be less than 1 might seem surprising, as it implies a phase velocity $\left(\frac{c}{n}\right)$ greater than c. However, this doesn't violate relativity, as signal velocity and group velocity are the relevant quantities for information transfer and are properly constrained.

What is the "Physical reason" for this behavior?

Page 15:

Continuing with the physical reason for anomalous dispersion in a gain medium:

"Kramers-Kronig relations link negative absorption (stimulated emission) to a dispersive phase shift." The Kramers-Kronig relations are profound mathematical relationships that connect the real and imaginary parts of the response function of a linear causal system. For electromagnetic waves interacting with a medium, the imaginary part of the susceptibility is related to absorption (or gain), and the real part is related to the refractive index (dispersion). So, if you have absorption or gain (which is negative absorption) over a certain frequency range, there must* be an associated variation in the refractive index over that same range and beyond. Stimulated emission, being negative absorption, fundamentally alters the dispersive properties of the medium compared to a normal absorbing medium.

* "The sign change in $\alpha(\nu)$ (gain instead of loss) forces a corresponding sign change in $\frac{dn}{d\nu}$." Here, $\alpha(\nu)$ is the absorption coefficient. If it's negative, we have gain. For an absorbing medium, you typically have normal dispersion around the absorption line. For a gain medium (negative $\alpha(\nu)$), this behavior of $n(\nu)$ flips, leading to the characteristic shape of anomalous dispersion we described – $n(\nu)$ decreasing then increasing as you sweep frequency through the gain line center. Specifically, $\frac{dn}{d\nu}$ is positive on the wings and negative in the center for an absorption line, and this pattern is modified for a gain line. The slide says "sign change in $\frac{dn}{d\nu}$," which means that across the resonance, the slope of $n(\nu)$ versus ν changes.

* This anomalous dispersion has a very practical consequence: "Resulting pulling of the cavity modes:" Because the refractive index $n(\nu)$ changes rapidly with frequency near the gain line center ν_0, and because the active cavity eigenfrequencies depend on $n(\nu)$ through $d^*(\nu)$, the actual lasing frequencies will be affected.

* Specifically, "Frequency of each mode is shifted towards ν_0 compared with the passive value $\nu_{q,\text{vac}}$." This is what we call "mode pulling." The gain medium's dispersive properties essentially "pull" the cavity resonant frequencies towards the center of the gain profile. If a passive cavity mode is slightly to one side of ν_0, the refractive index it experiences will shift it slightly closer to ν_0. This is a crucial effect for determining the precise emission frequency of a laser.

Page 16:

And to quantify this mode pulling effect:

* A "Quantitative expression derived in later slides (mode pulling formula)" will be provided. We won't derive it just yet, but we'll see a formula that explicitly calculates the active cavity frequency ν_a in terms of the passive cavity frequency $\nu_{q,vac}$ (which we can call ν_c or ν_r for cavity resonance), the gain line center frequency ν_0, and their respective linewidths. This formula will mathematically show how the lasing mode is a kind of "weighted average" of the cavity resonance and the gain line center.

Now, let's look at a visual representation of this anomalous dispersion.

Page 17:

This slide displays a graph crucial for understanding anomalous dispersion, titled "Slide 4: Anomalous Dispersion and Its Spectroscopic Importance," with a subtitle "Anomalous Dispersion for an Inverted Transition (within gain bandwidth)." Let's describe this graph carefully.

The horizontal axis represents frequency, ν (nu), with the center frequency of the inverted transition labeled as ν_0 (nu naught). Points like $\nu_0 - \Gamma$ (nu naught minus capital Gamma), $\nu_0 - \frac{\Gamma}{2}$, $\nu_0 + \frac{\Gamma}{2}$, and $\nu_0 + \Gamma$ are marked, suggesting Γ is related to the width of the transition. The region between approximately $\nu_0 - \Gamma$ and $\nu_0 + \Gamma$ is shaded light gray, labeled "Anomalous Dispersion (within gain bandwidth)".

There are two vertical axes and two curves.

The left vertical axis is "Absorption Coefficient ($\alpha(\nu)$)" – that's alpha of nu. The curve plotted against this axis is shown in red. Since this is an inverted transition (a gain medium), the absorption coefficient is negative. The red curve shows a negative peak (a dip, meaning maximum gain) centered at ν_0. It starts from a less negative value at low frequencies (e.g., $\nu_0 - \Gamma$), becomes most negative (maximum gain) at ν_0, and then rises back to less negative values at high frequencies (e.g., $\nu_0 + \Gamma$). For example, at ν_0, $\alpha(\nu)$ might be -1.0 in some arbitrary units shown, while at $\nu_0 \pm \Gamma$, it might be around -0.7. This red curve represents the gain profile of the laser medium.

The right vertical axis is "Refractive Index ($n(\nu)$)" – en of nu. The curve plotted against this axis is shown in blue. This is the curve illustrating anomalous dispersion. Far from resonance (e.g., to the left of $\nu_0 - \Gamma$ or to the right of $\nu_0 + \Gamma$), $n(\nu)$ is relatively flat. As we approach ν_0 from the low-frequency side, $n(\nu)$ decreases, dipping below a baseline value (which could be normalized to 1.0, as suggested by the dashed horizontal line at $n(\nu) = 1.00$). It reaches a minimum somewhere before ν_0 (around $\nu_0 - \frac{\Gamma}{2}$). Then, as frequency increases through ν_0, $n(\nu)$ rises sharply, crossing the baseline value ($n(\nu) = 1.00$ exactly at ν_0 for a symmetric gain profile), and then increases above the baseline, reaching a peak after ν_0 (around $\nu_0 + \frac{\Gamma}{2}$), before settling back towards the baseline at higher frequencies. The scale on the right shows values like 0.95, 1.00, and 1.05 for $n(\nu)$.

This graph perfectly visualizes what we discussed: 1. The red curve: gain is maximum at ν_0. $\alpha(\nu)$ is negative. 2. The blue curve: $n(\nu)$ shows the characteristic "S-shape" or "wiggle" through the resonance. On the low-frequency side of ν_0 ($\nu < \nu_0$), $n(\nu)$ can be less than 1 (if the gain is strong enough). On the high-frequency side ($\nu > \nu_0$), $n(\nu)$ is greater than 1. The slope $\frac{dn}{d\nu}$ is negative over a part of the range (e.g., between the peak and valley of $n(\nu)$), which is the hallmark of anomalous dispersion. Critically, at ν_0 itself, for a symmetric gain line, $n(\nu)$ is typically equal to its background value (often unity if the background medium is vacuum or if n is normalized), and the slope $\frac{dn}{d\nu}$ is at its most negative value here (if we consider the "anomalous" region centered around ν_0). However, looking at the blue curve, the steepest negative slope of $n(\nu)$ versus ν appears to be around ν_0. (Correction: the steepest negative slope $\frac{dn}{d\nu}$ is actually on the flanks for an absorption line. For a gain line as depicted, the slope $\frac{dn}{d\nu}$ is positive at ν_0. The "anomalous" behavior is the entire shape. The key feature for mode pulling is that $n(\nu)$ is different from its background value, and its derivative $\frac{dn}{d\nu}$ is significant). Let me re-examine the graph carefully: At ν_0, $n(\nu)$ is 1.00. To the left of ν_0, $n(\nu) < 1.00$. To the right, $n(\nu) > 1.00$. The slope $\frac{dn}{d\nu}$ at ν_0 is indeed positive and appears to be at its maximum positive value here. This is consistent with the Kramers-Kronig relations for a gain profile. The "anomalous" part refers to the entire dispersive feature linked to the resonance. The mode pulling depends on $n(\nu)$ itself, not just its derivative.

The key takeaway is that the refractive index $n(\nu)$ varies significantly across the gain bandwidth, and this variation directly impacts the cavity resonance conditions, leading to mode pulling.

Page 18:

Let's now precisely define the "Net Round-Trip Gain — Precise Definition." This is a critical quantity for determining if a laser will oscillate.

* We "Define $G(\nu, 2\,d)$ (capital Gee of nu comma two dee): amplitude (not intensity) gain factor for one round-trip of length $2\,d$ inside the resonator." This is very important: G is an *amplitude* gain factor. Intensity is proportional to amplitude squared. So, if the amplitude is multiplied by G after one round trip, the intensity will be multiplied by G^2. The arguments $(\nu, 2\,d)$ remind us that it's frequency-dependent and applies to a full round trip within the resonator of length 'd' (so round-trip path $2\,d$, or effective round trip path $2\,d^*$).

The mathematical definition provided is:

$$G(\nu, 2\,d) = \exp[-2\alpha(\nu)L - \gamma(\nu)]$$

Let's say this clearly: "Capital Gee of nu comma two dee equals the exponential of, open square bracket, minus two times alpha of nu times capital Ell, minus gamma of nu, close square bracket."

Now, let's break down the terms in the exponent: "where"

$\alpha(\nu)$ (alpha of nu): This is the "small-signal absorption coefficient." Its units are inverse meters (m^{-1}, or em to the minus one). Crucially, it is "negative if net gain." This means if the medium provides gain at frequency ν, $\alpha(\nu)$ will be a negative number. So, $-\alpha(\nu)$ would be positive. "Small-signal" implies this is the absorption/gain coefficient before* saturation effects become significant.

* L (capital Ell): This is the "length of the amplifying region" in meters (m, or em). This is the physical length of the gain medium that the light traverses.

The term $-2\alpha(\nu)L$ appears in the exponent. Why $2\alpha(\nu)L$? And why the minus sign with the alpha? If $\alpha(\nu)$ is the absorption coefficient, then for gain, $\alpha(\nu)$ is negative. Let's define $g(\nu) = -\alpha(\nu)$ as the gain coefficient (positive for gain). Then the term would be $\exp[2\,g(\nu)L - \gamma(\nu)]$. The slide uses $\alpha(\nu)$ as absorption coefficient, so if there's gain, $\alpha(\nu)$ itself is negative. Thus, $-\alpha(\nu)L$ represents the amplitude gain for a single pass through the medium of length L. For example, if light goes through, its amplitude E is $E_0 \cdot \exp[-\alpha(\nu)L]$. If $\alpha(\nu)$ is negative, say -0.1 per meter, and L is 1 meter, then $E = E_0 \cdot \exp(0.1)$. The amplitude increases. The factor of 2 in $-2\alpha(\nu)L$ will be explained on the next page, as it relates to the round trip.

Page 19:

Let's continue dissecting the terms in the round-trip amplitude gain factor $G(\nu, 2\,d) = \exp[-2\alpha(\nu)L - \gamma(\nu)]$.

- $\gamma(\nu)$ (gamma of nu): This is the "total passive loss factor per round-trip." It is "dimensionless" and represents all losses *other than* the absorptive/gain characteristic of the active medium's transition itself. It must be that $\gamma(\nu)$ is greater than 0 for there to be actual losses. This factor "Includes mirror transmission, scattering, diffraction, intracavity element losses." So, every time light completes a round trip, its amplitude is diminished by a factor of $\exp\left[-\frac{\gamma(\nu)}{2}\right]$ due to these passive losses on average per pass, or $\exp[-\gamma(\nu)]$ if gamma is defined for intensity loss per round trip. Here, since G is an amplitude gain factor and $\gamma(\nu)$ is in the exponent, $\gamma(\nu)$ here directly represents the amplitude attenuation exponent due to all these other passive losses over a round trip. A common form for this might be $\gamma(\nu) = -\ln(R_1 R_2 T_{\text{other}})$ where R's are mirror reflectivities and T_{other} includes other transmissive losses per round trip, all for amplitude. However, the way it's written suggests $\gamma(\nu)$ is an effective amplitude loss coefficient for the round trip. Let's assume $\gamma(\nu)$ is defined such that $\exp\big(-\gamma(\nu)\big)$ is the amplitude reduction factor from these passive losses per round trip.

- Now, the explanation for the term $-2\alpha(\nu)L$: "The first term $-2\alpha L$ appears twice because the light crosses the gain medium on the forward and backward paths." This clarifies that $-\alpha(\nu)L$ is the amplitude gain for a *single pass* through the gain medium of length L. So, for a round trip, the light passes through the medium twice (once forward, once backward if it's a linear cavity and the medium doesn't fill the whole cavity, or effectively traverses a length $2L$ if we consider the gain medium folding back on itself). Thus, the total amplitude gain from the active medium over one round trip is $\exp[-\alpha(\nu)L] * \exp[-\alpha(\nu)L] = \exp[-2\alpha(\nu)L]$. This makes perfect sense.

• Now we arrive at the "Threshold condition:" For laser oscillation to start and be sustained, the gain must exactly balance the losses. If the round-trip amplitude gain factor G is greater than 1, the amplitude grows with each pass. If it's less than 1, it decays. At threshold, it's exactly stable. So, the threshold condition is:

$$G(\nu_{\text{thr}}, 2\,d) = 1$$

"Capital Gee of nu sub thr (for threshold) comma two dee equals one." This means that at the threshold frequency ν_{thr}, the amplitude of the light wave remains unchanged after one complete round trip. The amplification from the gain medium precisely compensates for all the losses in the cavity. As the slide notes, this "means 'round-trip amplification equals all losses'." More precisely, the net round-trip amplitude change factor is unity. If $G = 1$, then the exponent must be zero: $-2\alpha(\nu_{\text{thr}})L - \gamma(\nu_{\text{thr}}) = 0$. This implies $-2\alpha(\nu_{\text{thr}})L = \gamma(\nu_{\text{thr}})$. Since for gain, α is negative, let's use $g(\nu) = -\alpha(\nu)$ as the gain coefficient (positive for gain). Then $2\,g(\nu_{\text{thr}})L = \gamma(\nu_{\text{thr}})$. This means the total amplitude gain from two passes through the medium, $2\,gL$, must equal the total amplitude loss exponent γ from all other sources per round trip.

Page 20:

Let's "Observe carefully" some properties of this round-trip gain G and the threshold condition.

* First, " G is frequency dependent through both $\alpha(\nu)$ and $\gamma(\nu)$." This is crucial. The gain coefficient $\alpha(\nu)$ (or $-\alpha(\nu)$) has its own spectral profile (e.g., Lorentzian, Gaussian). Also, losses $\gamma(\nu)$ can be frequency dependent. For example, mirror reflectivity might vary with frequency, or intracavity elements like etalons will have frequency-dependent transmission. Therefore, the overall round-trip gain G will vary with frequency, and laser oscillation will preferentially occur at frequencies where G is maximized and meets the threshold condition.

* Second, "The exponent is additive — changes in α or γ shift the logarithm of G linearly." Since $G = \exp[-2\alpha L - \gamma]$, if we take the natural logarithm, $\ln(G) = -2\alpha L - \gamma$. So, the logarithm of the round-trip amplitude gain is simply the sum of the gain contribution (which is $-2\alpha L$) and the loss contribution (which is $-\gamma$). If you change α (e.g., by changing the pump power) or if you change γ (e.g., by misaligning a mirror or inserting an absorbing filter), $\ln(G)$ changes linearly with these contributions. This additivity in the exponent (or logarithm of G) is often very convenient for analysis. For instance, at threshold, $\ln(G) = \ln(1) = 0$, so $-2\alpha L = \gamma$, as we noted earlier.

This understanding of G, its frequency dependence, and the threshold condition is fundamental to figuring out which frequencies will actually lase.

Page 21:

Now we explore "How a Probe Beam Reveals the Resonator Transmission Spectrum." This is a conceptual and often practical way to understand the resonant properties of a cavity, especially one containing an active medium (though it can be used for passive cavities too).

* "Thought experiment: send weak, broadband light with spectral intensity $I_0(\nu)$ into the cavity through mirror M_1."

Imagine we have our resonator (with or without an active medium inside). We take an external light source that emits over a wide range of frequencies (broadband). Its initial spectral intensity is $I_0(\nu)$ – that's power per unit area per unit frequency interval. We direct this light towards one of the cavity mirrors, say M_1, and some of it will be transmitted into the cavity. We keep the probe beam "weak" so that it doesn't significantly perturb the system, especially if there's an active medium that could be saturated.

* "Multiple internal reflections lead to constructive or destructive interference."

Once the light enters the cavity, it bounces back and forth between the mirrors. At each reflection, some light is transmitted out, and some is reflected back into the cavity. The light waves that have made different numbers of round trips inside the cavity will overlap and interfere.

At certain frequencies – the resonant frequencies – these multiple reflections will interfere constructively, leading to a buildup of light intensity inside the cavity and strong transmission through the output mirror.

At other frequencies (off-resonance), the interference will be destructive, leading to low intensity inside the cavity and weak transmission.

This is precisely the principle of a Fabry-Perot interferometer. By measuring the transmitted intensity as a function of frequency, we can map out the resonant modes of the cavity.

Page 22:

Let's quantify the transmission of this probe beam.

* "If amplification per pass is $G(\nu)$ (capital Gee of nu), mirror intensity transmission is T (capital Tee), reflectivity is R (capital Arr) (such that $R + T + \text{scattering} = 1$), then total transmitted intensity follows an Airy-type formula extended by gain:"

"This statement needs careful parsing. The $G(\nu)$ here seems to be defined differently from the round-trip amplitude gain on page 18. The context suggests $G(\nu)$ here might be the intensity* gain per pass through the active

medium, or something related to the overall gain within the cavity affecting the transmitted light. The Airy formula is typically for passive cavities. Let's look at the formula given:"

$$I_{\mathrm{T}}(\nu) = I_0(\nu) \cdot \frac{T^2\, G(\nu)}{[1 - G_{\mathrm{r}}(\nu)]^2 + 4\, G_{\mathrm{r}}(\nu) \sin^2\left(\frac{\phi}{2}\right)}$$

* "(I've used $G_{\mathrm{r}}(\nu)$ here for the term in the denominator based on typical Airy function form, where $G_{\mathrm{r}}(\nu)$ would be related to round trip reflectivity/gain. The slide writes $G(\nu)$ in the numerator and denominator, let's assume it's the same $G(\nu)$ for now and clarify.)"

* "Let's verbalize the formula as written on the slide carefully: 'Capital Eye sub Tee of nu equals Capital Eye sub zero of nu, times the fraction: Tee squared times Capital Gee of nu, all divided by, open square bracket one minus Capital Gee of nu close square bracket squared, plus four times Capital Gee of nu times sine squared of (phi divided by two).'"

Let's analyze this. $I_{\mathrm{T}}(\nu)$ is the transmitted spectral intensity, $I_0(\nu)$ is the incident spectral intensity. T is the intensity transmission coefficient of a single mirror (assuming both mirrors are identical for simplicity, or T is for the output mirror, and the formula may implicitly assume properties for the input mirror).

The $G(\nu)$ term here is tricky. If this formula is an extension of the standard Airy function for a Fabry-Perot with gain, $G(\nu)$ in the denominator is usually $\sqrt{R_1 R_2} \times G_{\text{medium}}$ where G_{medium} is the single-pass amplitude gain of the medium. And $G(\nu)$ in the numerator would be related to single pass gain.

However, if we assume $G(\nu)$ is the round-trip amplitude gain as defined on page 18 ($G = \exp[-2\alpha L - \gamma]$), this formula takes on a specific meaning related to regenerative amplification.

The term $[1 - G(\nu)]^2$ in the denominator is highly significant. If $G(\nu)$ approaches 1 (i.e., gain approaches losses, the threshold condition), this term approaches zero.

The $\sin^2\left(\frac{\phi}{2}\right)$ term dictates the resonant behavior. Transmission will be maximal when $\sin^2\left(\frac{\phi}{2}\right)$ is zero, meaning $\frac{\phi}{2}$ is an integer multiple of π (pi), or ϕ is an integer multiple of 2π.

* "with"

* "Phase advance per round-trip:" This is ϕ (the Greek letter phi).

Page 23:

The phase advance per round-trip, $\phi(\nu)$ (phi of nu), is given by:

$$\phi(\nu) = \frac{4\pi\, d^*(\nu)\, \nu}{c}$$

Let's understand this. The optical path for a round trip is \(2\,d^*(\nu)\). The wave number is* $k = \frac{2\pi}{\lambda} = \frac{2\pi\nu}{c}$. So, the phase accumulated over a path length P is $kP = \frac{2\pi\nu}{c}P$. For a round trip, \(P = 2\,d^(\nu)\), so \[\phi(\nu) = \frac{2\pi\nu}{c} \times 2\,d^*(\nu) = \frac{4\pi\nu d*^(\nu)}{c}.\] This is correct. This $\phi(\nu)$ is the total phase shift experienced by a wave in one full round trip inside the cavity.

* "Transmission maxima condition:" As we noted, transmission is maximized when the $\sin^2(\phi/2)$ term in the denominator of the Airy-like formula is zero. This occurs when $\phi/2 = q\pi$, where q is an integer. So, $\phi(\nu) = 2\pi q$. This is the condition for constructive interference and thus maximum transmission through the resonator.

* "Exactly the active-cavity eigenfrequency condition derived earlier." Let's check. Our phase is $\phi(\nu) = \frac{4\pi\nu d^*(\nu)}{c}$. Setting this equal to $2\pi q$:

$$\frac{4\pi\nu d^*(\nu)}{c} = 2\pi q$$

Dividing by 2π gives:

$$\frac{2\nu d^*(\nu)}{c} = q$$

Or, $2\,d^*(\nu) = q\left(\frac{c}{\nu}\right)$. This is precisely the active-cavity eigenfrequency condition

$$2\,d^*(\nu) = q\left(\frac{c}{\nu_{q,act}}\right)$$

that we found on page 12, if we identify ν here with $\nu_{q,act}$. So, yes, the maxima of the transmission spectrum of this "probed" cavity occur exactly at the resonant eigenfrequencies of the active cavity. This confirms that probing the cavity transmission is a way to find its resonant modes.

* Now, an "Important limit:" This relates back to the denominator of the transmission formula.

Page 24:

The important limit discussed here is critical for understanding the onset of laser oscillation:

* "As $G(\nu)$ approaches 1 from below (written as $G(\nu) \to 1^-$, meaning G is just less than 1) and ϕ (phi) satisfies the resonance condition ($\phi = 2\pi q$), the denominator tends to zero. Let's look at the denominator from page 22: $[1 - G(\nu)]^2 + 4\,G(\nu)\sin^2\left(\frac{\phi}{2}\right)$. If ϕ is at resonance, $\sin^2\left(\frac{\phi}{2}\right) = 0$, so the second term vanishes. The denominator becomes just $[1 - G(\nu)]^2$. If $G(\nu)$ is approaching 1 (meaning gain is approaching losses, we are nearing threshold), then $1 - G(\nu)$ is a small positive number, and $[1 - G(\nu)]^2$ is a very small positive number. So, the denominator indeed tends to zero.

* "This implies that $\frac{I_T}{I_0}$ diverges" (Capital Eye sub Tee divided by Capital Eye sub zero diverges). The ratio of transmitted intensity to incident intensity becomes extremely large. This divergence means that even for a very small input intensity I_0 (perhaps from spontaneous emission within the gain medium itself), we can get a very large transmitted (or output) intensity I_T.

* "⇒ laser oscillation set-up." (implies laser oscillation set-up). This is exactly what happens in a laser. When the round-trip amplitude gain $G(\nu)$ equals 1 (threshold), the cavity can build up a strong internal field and emit coherent radiation, even without an external probe beam, by amplifying the spontaneous emission. The formula essentially shows that the resonator's response becomes infinitely large at resonance when $G = 1$, signifying the system can self-sustain an oscillation.

Page 25:

This page shows a graph illustrating the "Resonator Transmission Spectrum: $\frac{I_T(\nu)}{I_0}$ vs. ν" (Capital Eye sub Tee of nu divided by Eye sub zero, versus ν). This plots the normalized transmitted intensity as a function of frequency, based on the Airy-type formula we just discussed.

The vertical axis is "Normalized Transmitted Intensity ($\frac{I_T(\nu)}{I_0}$)", ranging from 0.0 to 1.0. The horizontal axis is "Normalized Frequency ($\frac{\nu}{\nu_{FSR}}$)" (nu divided by nu sub Eff Ess Arr), centered around 1.00 (which would correspond to a resonance peak). It ranges from 0.90 to 1.10. Using ν_{FSR} here means we are looking at a single resonance peak, and the width of the x-axis covers about 20% of an FSR.

There are three curves plotted, for different values of the gain parameter G (which we assume is the $G(\nu)$ from the formula, representing round-trip amplitude gain): 1. Blue curve: $G = 0.8$. This represents a situation below threshold ($G < 1$). We see a resonance peak centered at the normalized frequency of 1.00. The peak height is relatively low, perhaps around 0.2 on the vertical axis, and the resonance is fairly broad. This is typical for a passive cavity with some losses, or an active cavity well below threshold.

2. Orange curve: $G = 0.99$. Here, the gain is very close to 1 (just below threshold). The resonance peak is still centered at 1.00, but it is now much taller (approaching 1.0, though the formula from page 22 with T^2G in numerator might not reach 1 unless T is also special) and significantly narrower. The peak intensity is dramatically increased, and the linewidth is greatly reduced as G approaches 1. This is a key feature: as you get closer to threshold, the cavity resonances become sharper and amplify more strongly. The orange curve on the graph appears to peak at about 0.8 perhaps.

3. Red curve: $G = 1.0$. This represents the system at the threshold of oscillation. According to our previous discussion, when $G = 1$ and we are at resonance, the denominator of the transmission formula becomes zero, and $\frac{I_T}{I_0}$ should diverge. The graph here shows the red curve for $G = 1.0$ as having a flat top at the maximum value of the y-axis (1.0) and being very narrow, almost like a spike that's been clipped at the top. This is a practical way of representing the onset of lasing – the transmission effectively becomes saturated or limited by other factors not in this simplified formula, or the output is now self-sustained laser light rather than transmitted probe light. The flat top at $\frac{I_T}{I_0} = 1$ suggests perhaps the intensity transmission T is also 1, or that G in the numerator is also 1. The key idea is the dramatic sharpening and increase in peak height as G approaches 1.

This graph beautifully illustrates the line-narrowing effect and the surge in transmitted power as a cavity with gain approaches the lasing threshold. It visually confirms our analysis from the previous pages.

Page 26:

Now we turn to a "Graphical Visualization of the Threshold Condition." This provides an intuitive way to see over what frequency range a laser might oscillate.

* "Procedure often used in laboratory alignment:"

1. "Plot the unsaturated gain curve $-2\alpha(\nu)L$ (negative ordinate implies amplification) vs frequency." Remember, $\alpha(\nu)$ is the absorption coefficient; it's negative for gain. So, $-\alpha(\nu)$ is the gain coefficient (let's call it $g_0(\nu)$ for small-signal gain coefficient). Thus, $-2\alpha(\nu)L$ is $2\,g_0(\nu)L$. This term represents the total logarithmic *amplitude* gain from two passes through the active medium of length L. "Unsaturated" means this is the gain profile of the medium *before* any significant laser intensity builds up and causes gain saturation. If plotted with positive values on the y-axis representing amplification, this curve will typically have a bell shape (e.g., Gaussian for Doppler broadening, Lorentzian for homogeneous broadening). The slide says "negative ordinate implies amplification" if plotting $-2\alpha(\nu)L$ directly with α negative, which could be confusing. It's clearer to think of $2\,g_0(\nu)L$ where g_0 is positive for gain. Let's assume the plot will show gain as positive.

2. "On the same axis plot the frequency-dependent total losses $\gamma(\nu)$." This $\gamma(\nu)$ is the total logarithmic amplitude loss per round-trip from all sources *other than* the specific lasing transition (e.g., mirror transmission, scattering, diffraction, absorption by other species or optics). This loss curve might be relatively flat if losses are not strongly frequency-dependent, or it might have its own spectral features.

3. "Subtract point-wise to obtain the net gain." This means, for each frequency ν, we calculate the difference between the gain $2\,g_0(\nu)L$ and the losses $\gamma(\nu)$.

Page 27:

Continuing with the graphical visualization:

The net gain, let's call it $\Delta\alpha(\nu)$ (Delta alpha of nu), though this notation might conflict if alpha is absorption, is defined as:

$$\Delta\alpha(\nu) = -2\alpha(\nu)L - \gamma(\nu)$$

"Delta alpha of nu equals minus two alpha of nu times Ell minus gamma of nu."

This $\Delta\alpha(\nu)$ is precisely the exponent in our round-trip amplitude gain $G(\nu) = \exp[\Delta\alpha(\nu)]$.

If $-2\alpha(\nu)L$ is plotted as gain (positive values), and $\gamma(\nu)$ as positive losses, then the net gain would be Gain - Loss. The formula $\Delta\alpha(\nu) = -2\alpha(\nu)L - \gamma(\nu)$ is directly the exponent of G. For G to be greater than 1 (net amplification), this exponent $\Delta\alpha(\nu)$ must be greater than 0.

4. "Those frequencies where $\Delta\alpha(\nu) > 0$ (Delta alpha of nu is greater than zero) satisfy $G > 1$ (Capital Gee is greater than one) and can lase."

This is the key. If the gain $(-2\alpha L)$ at a particular frequency exceeds the losses γ at that frequency, then $\Delta\alpha(\nu)$ is positive, meaning $G(\nu) = \exp(\text{positive number})$ will be greater than 1. This means that light at these frequencies will experience net amplification per round trip and can build up into laser oscillation. The range of frequencies where $\Delta\alpha(\nu) > 0$ is often called the "bandwidth for lasing."

* "The contact points $\Delta\alpha(\nu) = 0$ mark the threshold."

Where the net gain $\Delta\alpha(\nu)$ is exactly zero, the round-trip amplitude gain $G(\nu)$ is $\exp(0) = 1$. This is precisely the threshold condition. These are the frequencies at the very edges of the lasing band. At these points, gain just equals loss.

This graphical method allows one to quickly visualize if a laser is above threshold and over what range of frequencies it might lase, just by comparing the gain curve and the loss line.

Page 28:

This slide presents the "Graphical Visualization of Laser Threshold Condition" that we've just been discussing.

Let's analyze the graph.

The vertical axis is labeled "Gain / Loss Coefficient (Positive gain implies amplification)". Values range from -5 to 12 in arbitrary units.

The horizontal axis is "Frequency (ν)", with arbitrary units from 0 to 100.

There are three curves shown:

1. A solid blue line, labeled "$-2\alpha(\nu)L$ (Unsaturated Gain)". This is the gain curve. It's a bell-shaped curve, peaking at frequency $\nu = 50$ with a maximum gain value of about 10.5. It drops to near zero (or some small positive value representing the baseline) by $\nu = 20$ and $\nu = 80$. This represents the frequency-dependent amplification provided by the active medium over a round trip.

2. A dashed red horizontal line, labeled "$\gamma(\nu)$ (Total Losses)". This line is flat, at a constant value of approximately 4.0 across all frequencies. This represents a scenario where the passive cavity losses are constant with frequency (e.g., dominated by frequency-independent mirror transmission).

3. A solid green line, labeled "$\Delta\alpha(\nu)$ (Net Gain)". This curve is derived by subtracting the red loss line from the blue gain curve point-by-point $\Delta\alpha(\nu) = [-2\alpha(\nu)L] - \gamma(\nu)$. Where the blue gain curve is below the red loss line (e.g., for $\nu <\sim 33$ and $\nu >\sim 67$), the green net gain curve is negative. Where the blue gain curve is above the red loss line, the green net gain curve is positive. This positive net gain region is shaded light green and labeled "Lasing Region ($\Delta\alpha > 0$)". The green curve peaks where the difference between gain and loss is greatest, at $\nu = 50$, with a peak net gain of $10.5 - 4.0 = 6.5$.

The points where the green net gain curve $\Delta\alpha(\nu)$ crosses the zero axis (i.e., where $\Delta\alpha(\nu) = 0$) are marked as "Threshold." These occur where the blue gain curve intersects the red loss line, around $\nu = 33$ and $\nu = 67$. At these frequencies, gain exactly equals loss, $G = 1$.

Between these two threshold frequencies, $\Delta\alpha(\nu) > 0$, so $G > 1$, and the laser can oscillate within this frequency band. Outside this band, $\Delta\alpha(\nu) < 0$, $G < 1$, and oscillation is not possible because losses exceed gain.

This graph perfectly illustrates the concept: lasing occurs when the gain curve is above the loss line. The width of the lasing region depends on how much the peak gain exceeds the losses and the shape of the gain curve. If the peak of the blue curve were below the red line, there would be no lasing region at all.

Page 29:

Now we move to a "Case Study — Gas Laser With Doppler-Broadened Gain (Example 5.9a, Part 1)." This will allow us to apply the concepts we've learned to a concrete example, resembling a Helium-Neon (He-Ne) laser.

* "Given parameters (He-Ne-like):"

* First, "Line-centre small-signal absorption coefficient":

$$\alpha(\omega_0) = -0.01 \text{ cm}^{-1}$$

alpha at omega naught equals minus zero point zero one inverse centimeters. Notice this is given in terms of angular frequency ω_0 (omega naught) instead of ν_0 (nu naught), where $\omega = 2\pi\nu$. The units are cm^{-1}, which are CGS units. Since it's an absorption coefficient and we have gain, it's negative, as expected. This value $\alpha(\omega_0)$ is the absorption coefficient at the center of the gain line. The corresponding gain coefficient $g(\omega_0) = -\alpha(\omega_0) = +0.01 \text{ cm}^{-1}$.

* Next, "Gain cell length":

$$L = 10 \text{ cm}$$

Capital Ell equals ten centimeters. This is the physical length of the region containing the active gain medium.

Page 30:

Continuing with the parameters for our case study:

"Doppler FWHM (angular frequency)": This refers to the Full Width at Half Maximum of the Doppler-broadened gain profile, expressed in angular frequency units.

$$\delta\omega_{\text{D}} = 1.3 \times 10^9 \text{ Hz} \times 2\pi = 8.17 \times 10^9 \text{ rad s}^{-1}$$

"delta omega sub Dee equals one point three times ten to the ninth Hertz, times two pi, which equals eight point one seven times ten to the ninth radians per second."

Note that 1.3×10^9 Hz is 1.3 GHz, a typical Doppler width for He-Ne lasers. Multiplying by 2π converts this frequency width into an angular frequency width. So $\delta\nu_{\text{D}}$ (delta nu sub Dee) would be 1.3 GHz.

"Passive loss exponent per round-trip":

$$\gamma = 0.03$$

"gamma equals zero point zero three." This is the $\gamma(\nu)$ term from our threshold condition, assumed here to be constant (independent of frequency) and dimensionless, representing the logarithmic amplitude loss per round trip.

Now, we need the "Doppler-broadened Gaussian profile for amplification coefficient":

The term $-2\alpha(\omega)L$ is what we need for the exponent in the round-trip gain G. For a Gaussian profile, it's given by:

$$-2\alpha(\omega)L = -2\alpha(\omega_0)L \cdot \exp\left[-\left(\frac{\omega - \omega_0}{0.68\,\delta\omega_{\mathrm{D}}}\right)^2\right]$$

Let's verbalize this carefully: "minus two alpha of omega times Ell equals minus two alpha of omega naught times Ell, times the exponential of, open square bracket, minus, open parenthesis, fraction, omega minus omega naught, divided by, zero point six eight times delta omega sub Dee, close parenthesis, squared, close square bracket."

Let's analyze this expression:

$-2\alpha(\omega_0)L$ is the peak value of this gain term, occurring at the line center $\omega = \omega_0$.

The $\exp[\dots]$ term describes the Gaussian shape.

$\omega - \omega_0$ is the detuning from the line center.

The denominator $0.68 \cdot \delta\omega_{\mathrm{D}}$ in the Gaussian exponent is related to the standard deviation of the Gaussian. For a Gaussian $\exp\left(-\frac{x^2}{2\sigma^2}\right)$, the FWHM is $2\sqrt{2\ln 2}\,\sigma \approx 2.355\,\sigma$. If $\delta\omega_{\mathrm{D}}$ is the FWHM, then the relationship between $\delta\omega_{\mathrm{D}}$ and the term in the exponent might involve a conversion factor. Often, a Gaussian line shape is written as

$$g(\omega) = g(\omega_0)\exp\left[-\ln(2) \cdot \left(\frac{\omega - \omega_0}{\delta\omega_{\mathrm{D}}/2}\right)^2\right]$$

or

$$g(\omega) = g(\omega_0)\exp\left[-\left(\frac{\omega - \omega_0}{\Delta\omega}\right)^2\right]$$

where $\Delta\omega$ is related to the 1/e width.

The factor 0.68 seems specific;

$$\frac{1}{2\sqrt{\ln 2}} \approx 0.6005.$$

Or perhaps $\delta\omega_\text{D}/(2\sqrt{\ln 2})$ is used. Let's assume the 0.68 factor correctly relates $\delta\omega_\text{D}$ (FWHM) to the characteristic width in the exponent for this specific formulation. For a Gaussian $\exp(-x^2/w^2)$, FWHM is

$$2\sqrt{\ln 2}\, w.$$

So if $w = 0.68\, \delta\omega_\text{D}$, then FWHM would be

$$2\sqrt{\ln 2}\,(0.68\, \delta\omega_\text{D}) \approx 1.13\, \delta\omega_\text{D}.$$

This suggests $\delta\omega_\text{D}$ itself might not be the FWHM of *this exponential function*, but rather related to it. Or, the form $\exp[-(\text{argument})^2]$ means the function is half its max when

$$(\text{argument})^2 = \ln 2.$$

So,

$$\frac{\omega - \omega_0}{0.68\, \delta\omega_\text{D}} = \sqrt{\ln 2}$$

at half-max. Then

$$2(\omega - \omega_0) = \text{FWHM} = 2\sqrt{\ln 2}\,(0.68\, \delta\omega_\text{D}) \approx 1.133\, \delta\omega_\text{D}.$$

This suggests that $\delta\omega_\text{D}$ in the formula is actually a parameter approximately 0.88 times the FWHM of the gain curve. Let's assume the formula is as given and proceed. The numerical factor is important for precise calculations. Often, the Gaussian is written as

$$\exp\left[-4\ln(2)\left(\frac{\omega - \omega_0}{\delta\omega_\text{D}}\right)^2\right].$$

However, we'll use the formula as presented.

"Condition for positive net gain:"

As we saw, this is

$$-2\alpha(\omega)L - \gamma > 0,$$

or

$$-2\alpha(\omega)L > \gamma.$$

Page 31:

The condition for positive net gain, and thus for lasing to be possible, is:

$$-2\alpha(\omega)L > \gamma$$

"Minus two alpha of omega times Ell is greater than gamma."

Where $-2\alpha(\omega)L$ is the frequency-dependent round-trip logarithmic amplitude gain from the medium, and γ is the frequency-independent round-trip logarithmic amplitude loss.

The slide then instructs us to "Insert numbers ⇒ determine the angular frequency interval satisfying the inequality."

So, we need to plug in the given values:

$\alpha(\omega_0) = -0.01\,\text{cm}^{-1}$ $L = 10\,\text{cm}$

So, the peak gain term is

$$-2\alpha(\omega_0)L = -2 \times (-0.01\,\text{cm}^{-1}) \times (10\,\text{cm}) = -2 \times (-0.1) = +0.2$$

This is a dimensionless quantity, as it's an exponent.

The loss term is $\gamma = 0.03$.

The Doppler width parameter is $\delta\omega_\text{D} = 8.17 \times 10^9\,\text{rad s}^{-1}$.

So the inequality becomes:

$$0.2 \cdot \exp\left[-\left(\frac{\omega - \omega_0}{0.68 \cdot \delta\omega_\text{D}}\right)^2\right] > 0.03$$

We need to solve this for ω to find the range of frequencies where lasing can occur. This will give us the "lasing bandwidth."

Page 32:

We are now in "Part 2" of the "Case Study - Gas Laser With Doppler-Broadened Gain (Example 5.9a)."

* "Solve inequality numerically:"

The inequality we set up on the previous page was:

$$0.2 \cdot \exp\left[-\left(\frac{\omega - \omega_0}{0.68\,\delta\omega_\mathrm{D}}\right)^2\right] > 0.03$$

The slide actually writes:

$$0.02 \cdot \exp\left[-\left(\frac{\omega - \omega_0}{0.68\,\delta\omega_\mathrm{D}}\right)^2\right] > 0.03$$

Let me recheck my calculation of $-2\alpha(\omega_0)L$.

$\alpha(\omega_0) = -0.01\,\mathrm{cm}^{-1}$, $L = 10\,\mathrm{cm}$.

$-2 \cdot \alpha(\omega_0) \cdot L = -2 \cdot (-0.01\,\mathrm{cm}^{-1}) \cdot (10\,\mathrm{cm}) = -2 \cdot (-0.1) = +0.2$.

So the prefactor *should* be 0.2, not 0.02 as written on this slide for the numerical inequality.

If the slide intended $\alpha(\omega_0) = -0.001\,\mathrm{cm}^{-1}$, then $-2\alpha(\omega_0)L$ would be 0.02.

Let's assume there's a typo on this slide and the prefactor is indeed 0.2 from the previous inputs.

If we use 0.2:

$$0.2 \cdot \exp[-X^2] > 0.03,$$

where $X = \frac{\omega-\omega_0}{0.68\,\delta\omega_\mathrm{D}}$.

If we use 0.02 as on the slide:

$$0.02 \cdot \exp[-X^2] > 0.03.$$

Let's follow the slide's numbers for now and see the conclusion. If it uses 0.02:

1. "Rearrange:" Divide by 0.02:

$$\exp\left[-\left(\frac{\omega-\omega_0}{0.68\,\delta\omega_{\rm D}}\right)^2\right] > \frac{0.03}{0.02}$$

$$\exp\left[-\left(\frac{\omega-\omega_0}{0.68\,\delta\omega_{\rm D}}\right)^2\right] > 1.5$$

This is what's written on the slide.

So, the "0.02" prefactor was indeed used in the slide's subsequent steps. This implies that either $\alpha(\omega_0)$ was $-0.001\,\mathrm{cm}^{-1}$ or L was $1\,\mathrm{cm}$, or there's a factor of 10 error in $\alpha(\omega_0)$ or L values given on page 29, or in the calculation leading to 0.02.

Given $\alpha(\omega_0) = -0.01\,\mathrm{cm}^{-1}$ and $L = 10\,\mathrm{cm}$, the term $2\alpha(\omega_0)L$ is -0.2. So $-2\alpha(\omega_0)L$ is $+0.2$.

If the slide uses 0.02, it implies $g_0(\omega_0)L = 0.01$ for single pass gain, where $g_0 = -\alpha_0$.

So $2\,g_0(\omega_0)L = 0.02$.

Let's proceed by consistently using the slide's value of 0.02 as the starting point for the inequality on this page, while noting the discrepancy with page 29.

Page 33:

Continuing with the solution based on the rearranged inequality:

$$\exp\left[-\left(\frac{\omega-\omega_0}{0.68\,\delta\omega_{\rm D}}\right)^2\right] > 1.5$$

2. "Exponential on left cannot exceed 1 ⇒ the inequality has no solution unless the prefactor $-2\alpha(\omega_0)L$ is > 0.03. But the given magnitude $0.02 < 0.03$ contradicts that."

This is a critical point. The exponential term $\exp[-Y^2]$ where $Y^2 = \left(\frac{\omega-\omega_0}{0.68\,\delta\omega_{\mathrm{D}}}\right)^2$ is always positive. Since $Y^2 \geq 0$, $-Y^2 \leq 0$. Therefore, $\exp[-Y^2]$ must be less than or equal to $\exp(0)$, which is 1.

So,

$$\exp\left[-\left(\frac{\omega - \omega_0}{0.68\,\delta\omega_{\mathrm{D}}}\right)^2\right] \leq 1.$$

Our inequality from the previous page was

$$\exp\left[-\left(\frac{\omega - \omega_0}{0.68\,\delta\omega_{\mathrm{D}}}\right)^2\right] > 1.5.$$

Since the left side can never be greater than 1, it can certainly never be greater than 1.5.

This means that, *with the prefactor of* 0.02 for the gain term and a loss of 0.03, there is no frequency ω for which the inequality holds. The gain is insufficient to overcome the losses at any frequency.

The slide explains this: "Exponential on left cannot exceed 1 implies the inequality has no solution unless the prefactor $-2\alpha(\omega_0)L$ (which was 0.02 in the inequality setup) is greater than the loss γ (which was 0.03)."

This logic is slightly rephrased. The original inequality was

$$\text{PeakGainFactor} \times \text{GaussianShape} > \text{LossFactor}.$$

So,

$$0.02 \times \text{GaussianShape} > 0.03.$$

Since

$$\text{GaussianShape} \leq 1,$$

the maximum value of the left side is

$$0.02 \times 1 = 0.02.$$

So we are asking if

$$0.02 > 0.03,$$

which is false.

Therefore, the inequality

$$0.02 \times \exp\left[-\left(\frac{\omega - \omega_0}{0.68\,\delta\omega_{\mathrm{D}}}\right)^2\right] > 0.03$$

has no solution. The laser will not lase with these parameters (specifically, if the peak round-trip gain factor $-2\alpha(\omega_0)L$ is 0.02 and the loss γ is 0.03).

* "Interpretation:"

* "Pump power must be increased further until $-2\alpha(\omega_0)L = 0.03$." Or, more generally, until $-2\alpha(\omega_0)L > 0.03$. To reach threshold, the peak gain must at least equal the loss. The term $-2\alpha(\omega_0)L$ is proportional to the small-signal gain, which is in turn typically proportional to the pump power. So, increasing pump power increases this term. Lasing will only begin when the peak of the gain curve $-2\alpha(\omega_0)L$ at least touches the loss line $\gamma = 0.03$.

* "After threshold, the valid frequency range broadens as the pump increases because unsaturated $\alpha(\omega)$ becomes more negative."

Once $-2\alpha(\omega_0)L$ exceeds 0.03, there will be a range of frequencies around ω_0 where

$$-2\alpha(\omega)L > 0.03.$$

As the pump power increases further, $-2\alpha(\omega_0)L$ (the peak gain) increases, making the gain curve taller. This means the gain curve will exceed the loss line over a wider range of frequencies. So, the lasing bandwidth increases.

This case study, even with the apparent numerical inconsistency between slides, powerfully illustrates the threshold condition. If peak gain < loss, no lasing. If peak gain = loss, threshold at line center. If peak gain > loss, lasing over a finite bandwidth.

Page 34:

Continuing the interpretation, and perhaps implicitly correcting or clarifying the numerical situation:

* "Once $-2\alpha(\omega)L > 0.03$ at line centre, approximate stimulated-emission bandwidth:"

This phrasing implies that the condition for lasing is met when the gain at line center exceeds the loss of 0.03. (The slide uses ω here generically, but line center is ω_0.)

The text then gives a value for an "approximate stimulated-emission bandwidth": $\delta\omega \approx 2\pi \times 3\,\mathrm{GHz}$ "delta omega is approximately two pi times three gigahertz." This is an angular frequency bandwidth. The corresponding frequency bandwidth $\delta\nu$ would be $3\,\mathrm{GHz}$.

"as quoted in the text." This suggests this value comes from the textbook example this slide is based on (Example 5.9a). For this bandwidth to be achieved, the gain $-2\alpha(\omega_0)L$ must be sufficiently above the threshold $\gamma = 0.03$.

Let's see if we can work backwards. If the lasing bandwidth (FWHM of the region where gain > loss) is $\delta\omega_{\text{lase}}$, then at $\omega = \omega_0 \pm \frac{\delta\omega_{\text{lase}}}{2}$, we have $-2\alpha(\omega)L = \gamma$.

Using the Gaussian form:

$$-2\alpha(\omega_0)L \cdot \exp\left[-\left(\frac{\delta\omega_{\text{lase}}/2}{0.68 \cdot \delta\omega_{\text{D}}}\right)^2\right] = \gamma.$$

We have $\gamma = 0.03$. $\delta\omega_\text{D} = 8.17 \times 10^9\ \text{rad/s}$. $\delta\omega_\text{lase} = 2\pi \times 3 \times 10^9\ \text{rad/s}$.

So, $\frac{\delta\omega_\text{lase}}{2} = \pi \times 3 \times 10^9\ \text{rad/s} \approx 9.42 \times 10^9\ \text{rad/s}$.

The argument of the exp: $X_\text{half} = \frac{\pi\times3\times10^9}{0.68\times8.17\times10^9} \approx \frac{9.42}{0.68\times8.17} \approx \frac{9.42}{5.55} \approx 1.696$.

$X_\text{half}^2 \approx 2.876$.

So, $\exp[-2.876] \approx 0.0563$.

Then, $-2\alpha(\omega_0)L \cdot 0.0563 = 0.03$.

This would imply $-2\alpha(\omega_0)L = \frac{0.03}{0.0563} \approx 0.532$.

This peak gain value of 0.532 is much larger than the 0.2 we calculated from the initial parameters, or the 0.02 used in the slide's direct inequality.

This suggests that the 3 GHz bandwidth corresponds to a scenario where the pump power is significantly higher, leading to a much larger line-center gain.

This highlights that the actual lasing bandwidth is not just fixed but depends on how far above threshold the laser is operating. The initial parameters given on page 29 (leading to $-2\alpha(\omega_0)L = 0.2$) would result in a different, likely broader, lasing bandwidth than if it were 0.02, and certainly different from the one implied by the 3 GHz figure if that requires a gain of 0.532. The main lesson is the *process* of how to calculate it.

Page 35:

This page shows two graphs side-by-side, illustrating the "Gas Laser: Doppler-Broadened Gain vs. Loss," as a "Case Study based on Example 5.9a (He-Ne-like laser)." These graphs are interactive in principle, with sliders for "Peak Small-Signal Gain $G_0 = -2\alpha(\omega_0)L$ (Pump Level)".

Left Graph:

The slider for Peak Small-Signal Gain G_0 is set to 0.0195. This is very close to the 0.02 used in the problematic inequality calculation. The graph plots "Gain/Loss Coefficient" vs. "Frequency Detuning $\frac{\omega-\omega_0}{2\pi}$ (GHz)". The gain curve (blue, bell-shaped) is $G_0 \times$ GaussianShape. Its peak is at 0 detuning, and the height is 0.0195, labeled "$G_0 = 0.0195$". The loss line (red, dashed horizontal) is labeled "Loss $\gamma = 0.03$". Visually, the peak of the blue gain curve (0.0195) is clearly *below* the red loss line (0.03). In this scenario, Gain $<$ Loss everywhere. Therefore, no lasing is possible. This corresponds to the conclusion we reached when using the 0.02 prefactor: the laser is below threshold.

Right Graph:

Here, the slider for Peak Small-Signal Gain G_0 is set to 0.0565. The gain curve (blue, bell-shaped) now peaks at 0.0565, labeled "$G_0 = 0.0565$". The loss line (red, dashed horizontal) is still at "Loss $\gamma = 0.03$". Now, the peak of the blue gain curve (0.0565) is significantly *above* the red loss line (0.03). There is a region of frequencies around the center where the blue curve is above the red line. This region, where Gain $>$ Loss, is shaded in light blue, indicating the frequencies where lasing can occur. The width of this shaded region visually represents the lasing bandwidth. The detuning axis runs from about -2.5 GHz to $+2.5$ GHz. The shaded region seems to span from roughly -1.5 GHz to $+1.5$ GHz, making the FWHM of the lasing region approximately 3 GHz. This matches the "stimulated-emission bandwidth $\delta\nu \approx 3$ GHz" mentioned on the previous page. This confirms that the 3 GHz bandwidth occurs when the peak gain $-2\alpha(\omega_0)L$ is around 0.0565, not 0.2 or 0.02 or 0.532 (my earlier quick estimate for 0.532 was based on a FWHM definition for the lasing bandwidth, whereas this graphical method uses the full width where gain > loss). So, to get a 3 GHz lasing bandwidth (full width, not FWHM), we need $G_0 = 0.0565$ when $\gamma = 0.03$.

These two graphs perfectly illustrate the threshold behavior: 1. If peak gain G_0 (e.g., 0.0195) is less than loss γ (0.03), no lasing. 2. If peak gain G_0 (e.g., 0.0565) is greater than loss γ (0.03), lasing occurs over a finite bandwidth. Increasing G_0 further (by increasing pump power) would make the blue curve taller and the shaded lasing region wider.

PAGE 36:

Now we consider "Example 5.9a — Mode Counting in a 50 cm Resonator." We've determined the potential lasing bandwidth (e.g., $3\,\text{GHz}$ if G_0 is high enough). Now we need to see how many discrete cavity modes fit within this bandwidth.

* "Mirror separation (physical length) $d = 50\,\text{cm} = 0.5\,\text{m}$." "dee equals fifty centimeters, which equals zero point five meters." This defines our resonator.

* "Passive cavity FSR (Free Spectral Range):" We use the formula

$$\delta\nu = \frac{c}{2\,d}$$

$$\delta\nu = \frac{3.00 \times 10^8\,\text{m s}^{-1}}{2 \times 0.5\,\text{m}}$$

$$\delta\nu = \frac{3.00 \times 10^8\,\text{m s}^{-1}}{1.0\,\text{m}}$$

delta nu equals three point zero zero times ten to the eighth meters per second, divided by one point zero meters.

$$\delta\nu = 3.00 \times 10^8\,\text{Hz} = 300\,\text{MHz}$$

delta nu equals three point zero zero times ten to the eighth Hertz, which equals three hundred Megahertz. So, the axial modes of this passive $50\,\text{cm}$ cavity are spaced $300\,\text{MHz}$ apart.

* "Amplification bandwidth supporting oscillation ≈ 3 GHz." This is the result from the previous discussion (page 35, right graph) when the gain was sufficiently high $G_0 = 0.0565$, $\gamma = 0.03$. Let's call this $\Delta\nu_{\text{lase}}$ (Delta nu sub lase).

* Now, we want to find the "Number N_{axial} (En sub axial) of axial modes inside the gain band:" (More accurately, inside the bandwidth where gain exceeds loss).

Page 37:

To find the number of axial modes, N_{axial}, that can fit within the lasing bandwidth $\Delta\nu_{\text{lase}}$, we divide the total lasing bandwidth by the spacing between the modes (the FSR, $\delta\nu$):

$$N_{\text{axial}} = \frac{\Delta\nu_{\text{lase}}}{\delta\nu}$$

Using the values:

$$\Delta\nu_{\text{lase}} \approx 3\,\text{GHz} = 3000\,\text{MHz}$$

$$\delta\nu = 300\,\text{MHz}$$

So, $N_{\text{axial}} = \frac{3000\,\text{MHz}}{300\,\text{MHz}} = 10$

En sub axial equals three gigahertz divided by three hundred megahertz, which is approximately ten.

The slide shows the calculation:

$$N_{\text{axial}} = \frac{3\,\text{GHz}}{300\,\text{MHz}} \approx 10$$

* "These ten modes are potentially able to reach threshold; final number depends on saturation and competition effects elaborated in later slides."

This is a very important caveat. Just because 10 passive cavity modes fall within the region where small-signal gain exceeds loss, it doesn't automatically mean all 10 will lase simultaneously, or that they will lase with equal intensity.

Once lasing starts on one or more modes, the intensity builds up, and this leads to "gain saturation." The gain profile is no longer the "unsaturated" one we used for this calculation. Saturation can reduce the gain, potentially preventing some modes from reaching threshold or altering their power.

"Competition effects" also arise because different modes might draw energy from the same pool of excited atoms/molecules in the gain medium.

We will indeed elaborate on these complex but crucial effects in later slides. For now, we can say there are approximately 10 "candidate" modes.

Page 38:

Let's consider another case study: "Slide 11: Case Study — Solid-State / Dye Laser Broad Gain (Example 5.9b)." Solid-state lasers (like Ti:Sapphire) and dye lasers are known for having very broad gain bandwidths, much broader than typical gas lasers.

* "Typical dye laser gain FWHM (Full Width at Half Maximum):"

$$\Delta\nu_{\text{gain}} \approx 1 \times 10^{13}\,\text{Hz}$$

Delta nu sub gain is approximately one times ten to the thirteenth Hertz. This is 10 Terahertz (THz). This is an enormous bandwidth compared to the ~ GHz bandwidths for gas lasers.

* "Using same cavity length d = 50 cm:" As in the previous example, $d = 0.5$ m, which gives an FSR,

$$\delta\nu = \frac{c}{2\,d} = 300\,\text{MHz} = 3.0 \times 10^{8}\,\text{Hz}$$

Now, let's calculate the number of axial modes N_{axial} that would fit within this much broader gain bandwidth:

$$N_{\text{axial}} = \frac{\Delta\nu_{\text{gain}}}{\delta\nu}$$

$$N_{\text{axial}} = \frac{1 \times 10^{13}\ \text{Hz}}{3.0 \times 10^{8}\ \text{Hz}}$$

En sub axial equals one times ten to the thirteenth Hertz, divided by three point zero times ten to the eighth Hertz.

$$N_{\text{axial}} = \frac{1}{3} \times 10^{13-8} = \frac{1}{3} \times 10^{5} \approx 0.333 \times 10^{5} = 3.3 \times 10^{4}$$

The slide shows:

$$N_{\text{axial}} = \frac{1 \times 10^{13}\ \text{Hz}}{3.0 \times 10^{8}\ \text{Hz}} \approx 3 \times 10^{4}$$

approximately three times ten to the fourth. So, about thirty thousand axial modes!

* "Crucial observation:" This large number has significant implications.

Page 39:

The crucial observation stemming from the vast number of modes in broad-band media is:

* "In broad-band media, the passive cavity mode spacing is extremely fine compared with the gain window." We saw this: 300 MHz mode spacing versus a 10 THz ($10{,}000{,}000$ MHz) gain bandwidth. The modes are very densely packed under the gain curve. If the laser were to operate on all these modes, the output would be highly multimode and quasi-continuous in terms of its spectrum for many practical purposes, though still composed of discrete lines.

* This leads to a practical requirement for many applications: "Additional intra-cavity wavelength-selective elements (gratings, prisms, etalons) become essential for spectroscopy-grade single-mode operation." If you need a laser for high-resolution spectroscopy, you typically want a single,

well-defined, narrow frequency. With tens of thousands of potential modes, allowing the laser to choose freely would result in a broad, uncontrolled output.

Therefore, to achieve single-mode operation, or even to select a narrow band of frequencies, one must insert elements *inside* the laser cavity that introduce additional losses for unwanted frequencies, effectively narrowing the *net* gain profile.

Examples of such elements include:

* **Diffraction gratings:** These disperse light by angle according to wavelength, and can be used in Littrow or Littman-Metcalf configurations to select a narrow wavelength band for feedback into the cavity.

* **Prisms:** Similar to gratings, prisms disperse light by wavelength and can be used to tune the laser output.

* **Etalons:** These are themselves Fabry-Perot interferometers. A thin, solid etalon inserted in the cavity will have its own set of transmission peaks, but with a much larger Free Spectral Range than the main cavity (because the etalon is thin). By aligning a transmission peak of the intracavity etalon with a desired main cavity mode, and ensuring other main cavity modes fall on the lossy regions of the etalon's transmission curve, one can enforce single-mode operation. Often, a hierarchy of such selective elements is needed (e.g., a birefringent filter for coarse tuning, then one or more etalons for fine mode selection).

Without these elements, a dye laser or Ti:Sapphire laser would typically operate highly multimode.

Page 40:

This page displays a graph for the "Slide 11: Case Study — Solid-State / Dye Laser Broad Gain (Example 5.9b)," titled "Resonator Modes under Broad Dye Gain Curve."

The graph shows:

A very broad, light blue, bell-shaped curve representing the "Relative Gain" on the vertical axis (from 0.0 to 1.0, arbitrary units).

The horizontal axis is "Frequency (THz relative to gain center)," ranging from -5 THz to $+5$ THz. This means the gain curve shown has a Full Width at Half Maximum (FWHM) of roughly 10 THz, consistent with the $\Delta\nu_{\text{gain}} \approx 1 \times 10^{13}$ Hz given earlier.

Underneath this extremely broad gain curve, one would imagine the very densely packed cavity modes. They are so close together (300 MHz spacing) that they cannot be individually resolved on this frequency scale (which spans 10 THz $= 10{,}000$ GHz $= 10{,}000{,}000$ MHz). The diagram doesn't attempt to show individual modes but rather conveys that the entire continuous-looking gain profile is available.

Text annotations below the graph reiterate the parameters and results:

"Cavity Length $d = 50$ cm (Mode Spacing $\delta\nu = 300$ MHz)"

"Gain FWHM $\Delta\nu_{\text{gain}} \approx 10$ THz ($\approx 1 \times 10^{13}$ Hz)"

"Axial modes within FWHM: $N_{\text{axial}} \approx 3.3 \times 10^{4}$" (consistent with our 33,000 calculation)

"Total modes shown (gain > 0.5%): $\approx 4.0 \times 10^{4}$" This implies that if we consider all modes where gain is above a very small threshold (0.5% of peak gain), the number is even larger, around 40,000.

This visual powerfully emphasizes the challenge and necessity of mode selection techniques when working with such broad-gain-bandwidth laser media if a narrow or single frequency output is desired. Without selection, the laser would emit a very broad spectrum.

Page 41:

We now shift to discuss "Passive vs Active Cavity Linewidth and the Quality Factor." We've talked about resonant frequencies, but these resonances are not infinitely sharp; they have a certain linewidth.

* "Define passive resonator finesse F_{p} (Eff sub pee) and linewidth $\Delta\nu_{\mathrm{p}}$ (Delta nu sub pee) (half-power width). For mirrors of amplitude reflectivity r (lowercase arr) (where $T = 1 - r^2 -$ losses, with T being intensity transmission):"

* The "finesse," F_{p}, of a passive Fabry–Perot resonator is a measure of its sharpness or quality. Higher finesse means sharper resonances and a greater ability to resolve closely spaced spectral lines if used as an interferometer.

* The formula for finesse given here is:

$$F_{\mathrm{p}} = \frac{\pi\sqrt{r_{\mathrm{eff}}}}{1 - r_{\mathrm{eff}}}$$

(Corrected based on common finesse definitions, slide has $\pi\sqrt{r}/(1-r)$.)

* Actually, the most common formula for finesse in terms of reflectivity $R = r^2$ is

$$F_{\mathrm{p}} = \frac{\pi\sqrt{R}}{1 - R}$$

or

$$F_{\mathrm{p}} = \frac{\pi R^{1/2}}{1 - R}$$

If r is amplitude reflectivity, then $R = r^2$. So the formula

$$F_{\mathrm{p}} = \frac{\pi r}{1 - r^2}$$

is often used if r is high.

* The slide writes:

$$F_{\mathrm{p}} = \frac{\pi\sqrt{r}}{1 - r}$$

This seems unusual. Let's assume r here directly refers to the effective amplitude reflectivity that includes all losses, not just the mirror's intrinsic reflectivity. More standardly, finesse for high reflectivity mirrors is

$$F_{\mathrm{p}} \approx \frac{\pi}{1 - R_{\mathrm{eff}}}$$

where R_{eff} is effective intensity reflectivity, or

$$F_{\mathrm{p}} = \frac{\pi\sqrt{R_1 R_2}}{1 - \sqrt{R_1 R_2}}$$

for amplitude, assuming $r = \sqrt{R_1 R_2}$. Let's use the slide's formula: "Eff sub pee equals pi times square root of arr, divided by (one minus arr)." This r must be related to the round trip amplitude loss.

* Alternatively, if r is the amplitude reflectivity of each mirror, and losses are only due to transmission

$$T = 1 - r^2$$

(assuming no scattering/absorption in mirrors), then

$$F_{\mathrm{p}} = \frac{\pi\, r}{1 - r^2}.$$

* If r is the single mirror amplitude reflectivity, then $R = r^2$ is its intensity reflectivity. A common formula for finesse is

$$F = \frac{\pi\sqrt{R_1 R_2}}{1 - \sqrt{R_1 R_2} \times (1 - A)},$$

where A is single pass loss. Or simply

$$F = \frac{\text{FSR}}{\Delta\nu_{\text{p}}}.$$

* Let's assume the r in the slide's formula for F_{p} is an effective amplitude reflectivity per mirror accounting for all losses.

* The passive cavity linewidth $\Delta\nu_{\text{p}}$ (Delta nu sub pee), which is the FWHM of the resonance peak, is related to the Free Spectral Range (FSR, $\delta\nu$) and the finesse F_{p} by:

$$\Delta\nu_{\text{p}} = \frac{\delta\nu}{F_{\text{p}}}$$

"Delta nu sub pee equals delta nu divided by Eff sub pee." This is a fundamental definition: finesse is the ratio of the mode spacing to the mode linewidth. So, a high finesse cavity has very narrow resonance peaks compared to their separation.

* Now, what happens when we add gain? "When gain medium compensates losses and approaches threshold $G \to 1^{-}$ (Capital Gee approaches one from below), effective finesse rises:"

As the net round-trip amplitude gain G approaches 1, the resonances become much sharper, as we saw in the transmission spectrum graph (page 25). This means the effective *finesse of the* active* cavity increases dramatically.

Page 42

The effective finesse of the active cavity, let's call it F_{a}^{*}, is given by a formula analogous to the passive one, but incorporating the round- trip amplitude gain $G(\nu)$:

$$F_{\text{a}}^{*} = \frac{2\pi\sqrt{G(\nu)}}{1 - G(\nu)}$$

Notice the 2π here, whereas the passive finesse formula on the previous page had π. This factor of 2 difference can arise from defining finesse based on round- trip phase or other conventions. The key is the $\left(1 - G(\nu)\right)$ in the denominator. As $G(\nu)$ approaches 1, this denominator approaches zero, so F_a^* tends to infinity.

* Corresponding active-cavity linewidth:

The active cavity linewidth, $\Delta\nu_\text{a}$, is similarly related to the FSR ($\delta\nu$) and this active finesse F_a^*:

$$\Delta\nu_\text{a} = \frac{\delta\nu}{F_\text{a}^*}$$

Substituting the expression for F_a^*:

$$\Delta\nu_\text{a} = \delta\nu\left(\frac{1 - G(\nu)}{2\pi\sqrt{G(\nu)}}\right)$$

This boxed formula is extremely important. It shows how the linewidth of an active cavity mode behaves as the gain $G(\nu)$ approaches the threshold value of 1.

* Limit behaviour:

As $G(\nu) \to 1$, then $F_\text{a}^* \to \infty$ and $\Delta\nu_\text{a} \to 0$.

This is clear from the formulas. When $G(\nu) = 1$, the denominator $\left(1 - G(\nu)\right)$ becomes zero, so \(F_\text{a}^{}\) *diverges to infinity. Consequently, \(\Delta \nu_\text{a} = \frac{\delta \nu}{F_\text{a}^{}}*\) approaches zero.

This implies that, theoretically, at the threshold of laser oscillation, the linewidth of the lasing mode should become infinitely narrow! This is the famous line- narrowing effect in lasers.

Page 43:

However, the prediction of zero linewidth at threshold is an idealization.

* "In practice quantum noise sets a lower bound on the achievable laser linewidth (see Sect. 5.6 — Schawlow-Townes limit)."

This is a crucial point. Even if we could perfectly stabilize the cavity and operate exactly at $G = 1$ (which is also an idealization, as stable oscillation occurs slightly above threshold where gain saturation clamps G effectively to 1), there is a fundamental limit to how narrow the laser linewidth can be.

This limit arises from "quantum noise," specifically from the unavoidable spontaneous emission of photons into the lasing mode by the excited atoms/molecules in the gain medium. Each spontaneously emitted photon has a random phase relative to the coherent field already in the mode. These random phase contributions cause the phase of the laser field to undergo a random walk, which, in turn, leads to a finite frequency linewidth.

This fundamental quantum limit is known as the "Schawlow-Townes limit," named after Arthur Schawlow and Charles Townes, who first derived it in their seminal 1958 paper on the principles of lasers. We will discuss this in detail in Section 5.6.

So, while the classical theory predicts $\Delta\nu_{\text{a}} \to 0$, quantum mechanics imposes a non-zero floor on the laser linewidth. Achieving this Schawlow-Townes limit is a significant challenge and a benchmark for laser stability. Practical lasers often have linewidths much broader than this limit due to technical noise (vibrations, temperature fluctuations, pump power instability, etc.).

Page 44:

Now we consider "Off-Resonance Regions — Why Threshold Is Not Met" generally, even if the gain profile itself is broad.

* "Between successive resonances, phase ϕ deviates from $2\pi q$ by $\approx \pi$."

The resonant condition is $\phi = 2\pi q$. Exactly midway between two resonances, say for modes q and $q + 1$, the phase would be $\frac{2\pi q + 2\pi(q+1)}{2} =$

$2\pi q + \pi$. So, yes, the phase ϕ will differ from an integer multiple of 2π by about π in the regions far from any resonance (i.e., in the "valleys" between the transmission peaks of the Fabry-Perot).

* "For a Lorentzian resonance of half-width $\Delta\nu_{\mathrm{r}}$, loss factor $\beta(\nu)$ increases approximately ten-fold at $|\nu - \nu_0| \approx 3\Delta\nu_{\mathrm{r}}$."

This point refers to the shape of the cavity transmission peaks if they are Lorentzian. The loss factor mentioned here, $\beta(\nu)$, isn't the same γ we used for threshold, but rather related to how quickly the cavity's ability to support a mode deteriorates as we move off resonance.

A Lorentzian line shape falls off as

$$\frac{1}{1 + \left(\frac{\nu - \nu_0}{\Delta\nu_{\mathrm{r}}}\right)^2}.$$

At $\nu - \nu_0 = \Delta\nu_{\mathrm{r}}$ (at the half-width point), the function is $\frac{1}{2}$.

At $\nu - \nu_0 = 3\Delta\nu_{\mathrm{r}}$, the term $\left(\frac{\nu - \nu_0}{\Delta\nu_{\mathrm{r}}}\right)^2 = 3^2 = 9$. So the Lorentzian is $\frac{1}{1+9} = \frac{1}{10}$.

This means the transmission of the cavity (or its ability to sustain a mode) drops by a factor of 10 when you are about 3 half-widths away from the resonance center. This is a property of the Lorentzian lineshape. So, if a cavity mode is not excited, the light that is off-resonance from this mode experiences significantly higher effective losses or lower transmission.

* "Since the gain curve decays away from ν_0, net gain quickly turns negative outside the immediate vicinity of each mode."

This is the combined effect. We have a gain curve (e.g., Gaussian or Lorentzian) provided by the active medium. We also have the comb of cavity resonances, each with its own finite width (e.g., $\Delta\nu_{\mathrm{r}}$).

For lasing to occur, a cavity resonance peak must overlap significantly with a region where the gain medium provides amplification greater than the total cavity losses.

If a cavity mode is far from the center of the gain curve, the gain $g(\nu)$ available at that mode's frequency might be too low to overcome the losses γ.

Furthermore, even if a frequency falls under the broad gain envelope, if it's not very close to a cavity resonance, the effective "cavity loss" for that specific frequency is very high (due to destructive interference), so it won't lase. Lasing requires both sufficient gain AND being on a cavity resonance.

Page 45:

This leads to the conclusion regarding off-resonance regions:

- "Conclusion: lasing occurs only in narrow frequency slots centred at cavity resonances; broad 'valleys' remain non-lasing despite active medium present."

This summarizes the situation well. The laser does not simply emit light across the entire gain profile of the active medium. Instead, the output spectrum is "filtered" or "selected" by the optical resonator.

Lasing will only occur at, or very near, the discrete resonant frequencies of the cavity (the $\nu_{q,act}$).

And, of these cavity resonances, only those that also fall within the frequency region where the gain provided by the active medium is sufficient to overcome all losses will actually lase.

So, the laser output consists of a set of discrete frequencies, "narrow frequency slots."

The "broad 'valleys'" between these cavity resonances will not support lasing, even if those frequencies are well within the gain bandwidth of the active medium. This is because light at those frequencies does not satisfy the standing wave condition (or the $\phi = 2\pi q$ phase condition) and therefore does not build up coherently within the resonator. It experiences destructive interference.

This interplay between the continuous gain profile and the discrete cavity modes is what shapes the fundamental output spectrum of a laser.

Page 46:

Now we introduce another crucial concept: "Introduction to Gain Saturation — The Basic Concept." So far, we've mostly talked about the "small-signal" gain, which is the gain experienced by a very weak light field that doesn't significantly alter the populations of the energy levels in the gain medium. However, once laser oscillation starts and the intensity inside the cavity builds up, this is no longer true.

"Population inversion $\Delta N = N_2 - N_1$ created by pumping."

This is the prerequisite for gain. N_2 is the population density of the upper laser level, and N_1 is that of the lower laser level. Pumping (optical, electrical, etc.) establishes $N_2 > N_1$ for an inverted transition.

"Stimulated emission rate for one frequency component ν:"

The rate at which atoms/molecules are stimulated to emit photons from level 2 to level 1 by the presence of radiation at frequency ν is given by:

$$R_{\text{stim}}(\nu) = B_{21} \cdot \rho(\nu) \cdot g(\nu - \nu_0) \cdot N_2$$

Let's define these terms:

B_{21}: This is the "Einstein B-coefficient" for stimulated emission. It's a fundamental property of the specific atomic or molecular transition, quantifying how strongly the transition interacts with light. Its units would be related to $(\text{energy density})^{-1}\,\text{time}^{-1}$ or $(\text{intensity})^{-1}\,\text{frequency}\,\text{time}^{-1}$. Standard units are $\text{m}^{3\,\text}{J}^{-1}\,\text{s}^{-2}$.

Page 47:

Continuing with the terms in the stimulated emission rate $R_{\text{stim}}(\nu) = B_{21} \cdot \rho(\nu) \cdot g(\nu - \nu_0) \cdot N_2$:

* $\rho(\nu)$ (rho of nu): This is the "spectral energy density" of the radiation field at frequency ν, inside the gain medium. Its units are Joules per cubic meter per Hertz $\left(\text{J m}^{-3}\ \text{Hz}^{-1}\right)$. It's a measure of how much electromagnetic energy is present at that frequency.

* $g(\nu - \nu_0)$ (gee of nu minus nu naught): This is the "normalized line-shape function." It describes the frequency profile of the transition (e.g., Lorentzian, Gaussian), centered at ν_0. It's normalized such that its integral over all frequencies is unity $\left(\int g(\nu - \nu_0)\, d\nu = 1\right)$. It ensures that stimulated emission is most efficient at the line center and falls off according to the transition's natural shape. It's dimensionless.

* N_2 (En sub two): This is the population density of the upper laser level (number of atoms/molecules per unit volume in state 2).

Now, the crucial consequence of stimulated emission:

* "As intra-cavity intensity I (capital Eye) grows, $\rho(\nu) \propto I$ (rho of nu is proportional to Eye) grows, hence R_{stim} depletes N_2."

The spectral energy density $\rho(\nu)$ is directly proportional to the light intensity $I(\nu)$ inside the cavity. As the laser starts to oscillate, $I(\nu)$ increases. This increases $\rho(\nu)$, which in turn increases the stimulated emission rate R_{stim}.

Each stimulated emission event takes one atom/molecule from the upper state N_2 to the lower state N_1. So, a high R_{stim} leads to a rapid depletion of the population N_2 (and an increase in N_1). This reduces the population inversion $\Delta N = N_2 - N_1$.

* "Under steady-state operation:"

In continuous-wave (CW) laser operation, a steady state is reached where the rate at which the pump excites atoms into N_2 is balanced by all the rates depleting N_2 (stimulated emission, spontaneous emission, non-radiative decay). The dominant depletion process for N_2 in a lasing mode is stimulated emission. So, "pump rate = stimulated emission rate" (approximately, for the lasing transition).

* This leads to the population inversion being "pinned" to a smaller value: $\Rightarrow$ inversion is 'pinned' to a smaller value ΔN_{thr} (Delta En sub thr).

As intensity rises, stimulated emission rises, which reduces the inversion. The gain is proportional to the inversion. The system will stabilize when the gain (which depends on the now reduced inversion) is just enough to balance the losses. This means the inversion ΔN will be clamped or "pinned" at the value it needs to have to sustain oscillation at that particular intensity, which is precisely the threshold inversion ΔN_{thr} required for gain to equal loss. If the pump tries to increase the inversion above this, the intensity will rise, increasing R_{stim}, which will drive the inversion back down to ΔN_{thr}. This is a self-regulating negative feedback mechanism.

Page 48:

And the name for this crucial process is:

* "This process is termed gain saturation." The gain of the medium becomes saturated because the population inversion is depleted by the strong stimulated emission driven by the laser's own internal field. The gain coefficient $g(\nu)$ is no longer the "small-signal" gain $g_0(\nu)$ (which depends on the unsaturated inversion), but rather a reduced, "saturated" gain $g_{\text{s}}(\nu)$ which depends on the (now lower) inversion ΔN_{thr}.

* "It introduces nonlinearity necessary for amplitude stability." This is a profound point. Gain saturation is a non-linear process because the gain itself becomes dependent on the intensity of the light.

If there were no gain saturation (i.e., if gain were independent of intensity), then once gain exceeded loss, the laser intensity would theoretically grow indefinitely, which is unphysical.

Gain saturation provides the mechanism for the laser to reach a stable output power. If the intensity tries to rise above the steady-state value, the gain saturates more, reducing the net gain, and pushing the intensity back down. If the intensity drops, gain saturation lessens, net gain increases, and intensity is pushed back up. This is what allows a laser to operate at a constant output power once it's above threshold.

So, gain saturation is not a detrimental effect to be avoided; it's an essential physical process for the stable operation of virtually all lasers.

Page 49:

Now we look at how gain saturation manifests in different types of broadened lines, starting with "Homogeneous Broadening — Uniform Saturation Across Frequency."

* "For a homogeneously broadened line:"

What is homogeneous broadening? In a homogeneously broadened transition, all atoms or molecules in the gain medium behave identically. They all have the same line center frequency and the same lineshape (e.g., a natural lifetime broadened line, or a collisionally broadened line in a gas at high enough pressure, or transitions in perfect crystals at low temperatures). There's no statistical distribution of resonance frequencies among the atoms.

* The consequence is: "All atoms/molecules 'see' every frequency component equally well."

More precisely, every atom has the same response curve. So, if there is a light field at a particular frequency ν within the gain profile, it can interact with *any* of the atoms in the upper laser state because they all share that same gain profile.

* Therefore, "Saturation of one laser mode reduces gain for the entire line profile."

This is the key characteristic of homogeneous saturation. If a single laser mode at frequency ν_{L} starts to lase and builds up intensity, it will deplete the population inversion N_2. Since all atoms contribute to the entire gain profile, this depletion of N_2 reduces the gain *across the whole profile*, not just at ν_{L}. The gain curve $g(\nu)$ retains its shape but is uniformly scaled down in amplitude.

* "Standard saturation formula (derivation recalled from rate equations):"

The effect of this saturation on the gain coefficient can be quantified. We usually derive this from a steady-state analysis of the rate equations for the populations N_1 and N_2.

Page 50:

The standard saturation formula for a homogeneously broadened gain coefficient, $\alpha_{s,\mathrm{hom}}(\nu)$ (alpha sub ess comma hom of nu), where 's' denotes saturated and 'hom' for homogeneous, is:

$$\alpha_{s,\mathrm{hom}}(\nu) = \frac{\alpha_0(\nu)}{1 + S}$$

which can also be written as:

$$\alpha_{s,\mathrm{hom}}(\nu) = \frac{\alpha_0(\nu)}{1 + \frac{I}{I_{\mathrm{s}}}}$$

alpha sub ess comma hom of nu, equals alpha sub zero of nu, divided by, open parenthesis one plus capital Ess close parenthesis. This also equals alpha sub zero of nu, divided by, open parenthesis one plus capital Eye over capital Eye sub ess close parenthesis.

Let's define these terms: "where"

* $\alpha_0(\nu)$ (alpha sub zero of nu): This is the "small-signal gain coefficient (before saturation)." (Note: if alpha is absorption coefficient, then gain is $-\alpha$. If this formula is for the gain coefficient itself, then $\alpha_0(\nu)$ should be positive for gain. Let's assume $\alpha_0(\nu)$ here represents the magnitude of the gain, or that it is the absorption coefficient which becomes less negative, i.e. closer to zero or even positive if saturation is very strong and turns gain into loss). More commonly, $g_{\text{s}}(\nu) = \frac{g_0(\nu)}{1+\frac{I}{I_{\text{s}}}}$ where g is gain. If α is absorption, then $\alpha_{\text{s}} = \frac{\alpha_0}{1+S}$ would mean a negative α_0 (small signal gain) becomes less negative. This is consistent. So $\alpha_0(\nu)$ is the unsaturated absorption coefficient (negative for gain).

* I_{s} (capital Eye sub ess): This is the "saturation intensity." It is a characteristic parameter of the gain medium and the transition. It is defined as "the intensity for which $S = 1$," meaning the intensity at which the gain coefficient is reduced to half its small-signal value (i.e., $\alpha_{s,\text{hom}}(\nu) = \frac{\alpha_0(\nu)}{2}$). Its units are typically Watts per square meter (W/m²) or Watts per square centimeter (W/cm²).

* $S \equiv \frac{I}{I_{\text{s}}}$ (Capital Ess is identically equal to Capital Eye over Eye sub ess): This is the "dimensionless saturation parameter." It's the ratio of the actual intensity I of the light at frequency ν to the saturation intensity I_{s}. S quantifies how strongly the transition is being saturated. If $S \ll 1$, saturation is weak. If $S \gg 1$, saturation is strong.

* "Consequence:"

Page 51:

The consequence of this uniform saturation in a homogeneously broadened medium is:

* "Once one mode reaches high intensity, the gain for neighbouring modes drops as well ⇒ mode competition, often resulting in single-frequency output for perfectly homogeneous media."

This is a very significant outcome. Imagine several cavity modes lie under the homogeneously broadened gain curve. Let's say the mode closest to the peak of $\alpha_0(\nu)$ starts to lase first.

As its intensity I builds up, it saturates the entire gain profile $\alpha_0(\nu)$, reducing it to

$$\alpha_s(\nu) = \frac{\alpha_0(\nu)}{1 + \frac{I}{I_s}}.$$

This reduction in gain occurs *for all frequencies*, including those of the other nearby cavity modes. If the first mode is strong enough, it can suppress the gain for all other modes below the threshold level (where gain equals loss).

This is "mode competition": the modes compete for the available saturated gain. In a perfectly homogeneous medium, the mode that is best positioned (usually highest gain-to-loss ratio) will tend to capture all the available energy and suppress all other modes.

This "often results in single-frequency output." This is why lasers based on homogeneously broadened gain media (like some solid-state lasers or dye lasers when carefully designed) have a natural tendency towards single-mode operation, provided other effects like spatial hole burning (which we'll see next) are minimized. If one mode can use up all the inversion, no other mode can start.

This is a powerful mechanism for achieving spectrally pure laser output.

Page 52:

Now we contrast this with the case of "Inhomogeneous Broadening — Selective Saturation and Hole Burning."

"In inhomogeneous ensembles (Doppler, Stark, inhomogeneous crystal fields):"

What is inhomogeneous broadening? In this case, different atoms or molecules in the gain medium have slightly different resonant frequencies. They are not all identical. This can arise from several physical mechanisms:

◦ **Doppler broadening:** In a gas, atoms/molecules are moving with a distribution of velocities. Due to the Doppler effect, each atom "sees" the light (or emits light) at a frequency shifted according to its velocity component along the light propagation direction. The overall gain profile is a superposition of many narrower, Doppler-shifted homogeneous lines.

◦ **Stark broadening:** Electric fields (e.g., from ions and electrons in a plasma) can shift energy levels, and if these fields are non-uniform, they lead to a distribution of transition frequencies.

◦ **Inhomogeneous crystal fields:** In some solid-state gain media (especially glasses or imperfect crystals), individual active ions may sit in slightly different local environments within the host material. These variations in the local crystal field can lead to slightly different transition frequencies for different ions.

The overall observed gain profile is the envelope of these many distinct, narrower sub-profiles.

"Individual molecules respond only to frequencies within their narrow homogeneous sub-profile $\Delta\nu_{\text{hom}}$ (Delta nu sub hom)."

Each atom or molecule (or group of atoms with the same local conditions/velocity) has its own intrinsic, much narrower homogeneous lineshape (e.g., determined by natural lifetime or collisions), centered at its specific resonant frequency. It only interacts strongly with light that falls within this narrow homogeneous width $\Delta\nu_{\text{hom}}$.

"A strong monochromatic field at ν_{L} (nu sub Ell) 'burns a hole' in the overall population distribution $\Delta N(\nu)$ (Delta En of nu)."

This is the critical difference from homogeneous saturation. If you apply a strong laser field at a specific frequency ν_{L}, it will primarily interact with and saturate only those atoms/molecules whose individual resonant frequencies are close to ν_{L} (i.e., within about $\Delta\nu_{\text{hom}}$ of ν_{L}). These atoms will have their population inversion $N_2 - N_1$ depleted. Atoms whose resonant frequencies are far from ν_{L} will be largely unaffected.

This creates a "hole" in the distribution of population inversion when plotted against frequency. At ν_{L}, the inversion is reduced (saturated), while away from ν_{L}, the inversion remains largely at its unsaturated value. This is known as "spectral hole burning." The "hole" appears in ΔN as a function of ν, where ν here represents the resonant frequency of a particular subgroup of atoms.

Page 53:

Let's look at the "Calculated saturated gain coefficient" for an inhomogeneously broadened line, $\alpha_{s,inh}(\nu)$. The formula is different from the homogeneous case:

$$\alpha_{s,inh}(\nu) = \frac{\alpha_0(\nu)}{\sqrt{1+S}}$$

This can also be written as:

$$\alpha_{s,inh}(\nu) = \frac{\alpha_0(\nu)}{\sqrt{1+\frac{I}{I_{\text{s}}}}}$$

Notice the square root in the denominator, $\sqrt{1+S}$, compared to $(1+S)$ for the homogeneous case. This means that for a given saturation parameter S

(or $\frac{I}{I_s}$), an inhomogeneously broadened line saturates more slowly, or less severely, than a homogeneously broadened one. The gain reduction is less pronounced.

* "Two competing trends with rising I (intensity):" Why this $\sqrt{1+S}$ dependence? It arises from two effects:

1. "Population depletion ⇒ gain $\propto \frac{1}{1+S}$." For the specific subgroup of atoms that are resonant with the saturating field at ν_L, their population inversion *is* depleted according to the homogeneous saturation law, so their contribution to gain reduces as $\frac{1}{1+S}$. This is the "depth" of the hole burned.

2. "Spectral hole broadening ⇒ effective number of contributing molecules $\propto \sqrt{1+S}$." As the intensity I of the saturating field increases, it doesn't just saturate the atoms exactly at ν_L more deeply. It also starts to saturate atoms whose resonant frequencies are slightly further away from ν_L. The "hole" burned in the spectral distribution of inversion becomes not only deeper but also wider. This effect is known as power broadening of the homogeneous linewidth. The width of the hole, $\Delta\nu_{hole}$, increases with intensity, approximately as

$$\Delta\nu_{hom} \cdot \sqrt{1+S}$$

The "effective number of contributing molecules" to the gain at frequency ν_L is related to this hole width. As the hole broadens, more atoms from the inhomogeneous distribution are brought into resonance with the saturating field (due to power broadening of their individual homogeneous lines), but they are also saturated.

The "Net result given above" $\frac{1}{\sqrt{1+S}}$ comes from a more detailed calculation that combines these effects. Essentially, while the peak of the hole saturates like $\frac{1}{1+S}$, the hole also broadens, bringing more atoms into play. The overall effect on the macroscopic gain coefficient $\alpha_{s,inh}(\nu)$ at the frequency of the saturating light results in the $\frac{1}{\sqrt{1+S}}$ dependence. This is typically derived for a Lorentzian homogeneous lineshape and a much broader inhomogeneous profile.

Page 54:

This page and the next visually illustrate "Spectral Hole Burning in Inhomogeneously Broadened Media." We have two plots here, (a) and (b). The vertical axis for both is "Population Inversion (ΔN)" (Delta En), ranging from 0.0 to 1.0 (arbitrary units). The horizontal axis is "Frequency (ν)", ranging from -100 to 100 (arbitrary units), centered at ν_0 (nu naught), which is the center of the inhomogeneous profile.

(a) "Unsaturated Inhomogeneous Profile ($\Delta N_0(\nu)$)" (Delta En sub naught of nu)

This shows a broad, smooth, typically Gaussian-shaped curve representing the distribution of population inversion before any strong monochromatic light is applied. The peak is at ν_0. This is the overall profile that determines the small-signal gain $\alpha_0(\nu)$.

(b) "Narrow Hole (Low Intensity)"

Now, a monochromatic laser field at frequency ν_L (nu sub Ell) is applied. ν_L is shown as a vertical dashed line, slightly offset from ν_0 for illustration (though the hole would be centered at ν_L).

The population inversion profile is now modified. We see a "dip" or a "hole" burned into the profile, centered at ν_L.

The annotation "$S = 1.0$" indicates that the intensity of the laser field I_L is equal to the saturation intensity I_s for the individual homogeneous packets.

At ν_L, the population inversion is reduced significantly (by a factor of 2 for $S = 1$ if it were purely homogeneous saturation at the peak of the hole).

Away from ν_L, the profile remains largely unperturbed, still following the original unsaturated shape.

The width of this hole is related to the homogeneous linewidth $\Delta\nu_{hom}$, possibly slightly power-broadened if $S = 1.0$ leads to some broadening.

This graph clearly shows that only a selective group of atoms (those resonant with ν_L) has been saturated. Atoms with resonant frequencies far from ν_L still have their full inversion and can contribute to gain at those other frequencies.

Page 55:

This is panel (c) of the spectral hole burning illustration, titled "(c) Wider, Deeper Hole (High Intensity)." The axes are the same: "Population Inversion (ΔN)" vs. "Frequency (ν)". The monochromatic laser is still at ν_L (dashed vertical line).

Now, the annotation says $S = 10.0$. This means the intensity I_L is ten times the saturation intensity I_s. This is a strong saturation regime. What do we observe in the population inversion profile? The hole centered at ν_L is now much "deeper" – the population inversion at ν_L is very strongly reduced, much closer to zero.

And, crucially, the hole is also significantly "wider" than in the $S = 1.0$ case. This is due to power broadening. The strong field at ν_L can now effectively saturate atoms whose resonant frequencies are further away from ν_L because their individual homogeneous lines have been broadened by the strong field. The formula for the power-broadened homogeneous linewidth is approximately $\Delta\nu_{hom}' = \Delta\nu_{hom} \times \sqrt{1+S}$.

With $S = 10$, $\sqrt{1+10} = \sqrt{11} \approx 3.3$. So the hole width is roughly 3.3 times the low-intensity homogeneous linewidth.

These three panels (unsaturated, $S = 1$ hole, $S = 10$ hole) beautifully demonstrate: 1. Selective saturation at the laser frequency ν_L. 2. The hole becomes deeper as intensity increases. 3. The hole becomes wider (power broadening) as intensity increases.

The key implication is that even if one mode burns a hole, other modes at different frequencies might still be able to lase if they can find enough unsaturated atoms elsewhere in the inhomogeneous profile. This makes simultaneous multi-mode oscillation more likely in inhomogeneously broadened lasers compared to homogeneously broadened ones.

Page 56:

We now turn to a different kind of hole burning, not in the spectral domain, but in the spatial domain: "Spatial Hole Burning — Standing-Wave Saturation Pattern." This occurs specifically in linear (standing-wave) resonators.

* "In a linear resonator the field is a standing wave:" A linear resonator is typically formed by two mirrors. The electromagnetic field inside, for a given lasing mode, forms a standing wave pattern. This standing wave has fixed positions of nodes (where the electric field amplitude is always zero) and antinodes (where the electric field amplitude oscillates with maximum strength).

The electric field $E(z)$ along the cavity axis 'z' can be described by:

$$E(z) = E_0 \cos\left(\frac{2\pi z}{\lambda}\right)$$

or

$$E(z) = E_0 \sin\left(\frac{2\pi z}{\lambda}\right)$$

depending on boundary conditions.

"Eee of zee equals Eee sub zero times cosine of (two pi zee over lambda)." (Assuming a mirror at $z = 0$ means E must be zero, so a sine function might be more appropriate if $z = 0$ is a mirror. However, intensity is E^2, so $\cos^2$ or $\sin^2$ behave similarly in terms of periodic zeros).

The important part is the sinusoidal spatial variation.

This results in "nodes at mirrors, antinodes inside." (More accurately, nodes at perfectly conducting mirrors. Antinodes occur at positions where $\frac{2\pi z}{\lambda}$ is an integer multiple of π for cosine, or $\left(n + \frac{1}{2}\right)\pi$ for sine relative to a node).

* "Local intensity I(z) ∝ |E(z)|² = E_0 ² cos²(2πz/λ)." "Capital Eye of zee is proportional to the magnitude of Eee of zee squared, which equals Eee sub zero squared times cosine squared of (two pi zee over lambda)."

The intensity of the light is proportional to the square of the electric field amplitude. So, the intensity $I(z)$ also has a spatially varying pattern:

It will be maximum at the antinodes of the E-field.

It will be zero at the nodes of the E-field.

This creates a periodic grating-like pattern of high and low intensity along the cavity axis.

* "Saturation therefore occurs only at the antinodes." Or, more precisely, saturation is strongest at the antinodes where the intensity is highest, and weakest (or non-existent) at the nodes where the intensity is lowest (or zero). So, the gain medium experiences spatially non-uniform saturation. Atoms located near the antinodes of the standing wave will be strongly saturated. Atoms located near the nodes will experience very little intensity and thus remain largely unsaturated.

Page 57:

Consequences of this spatial hole burning:

* "Nodes remain unsaturated ⇒ residual inversion." At the spatial locations of the standing wave nodes, the intensity is zero (or very low). Therefore, the gain medium at these locations is not significantly saturated. There remains a substantial population inversion (residual inversion) in these regions.

* "Another mode shifted by $\lambda/4$ can exploit these unsaturated regions ⇒ multiple modes may co-exist even in homogeneous media." This is a critical consequence for multi-mode operation. Consider a lasing mode with wavelength λ, which creates a standing wave and burns spatial holes. Now, think about another potential cavity mode. If this second mode has a slightly different wavelength (and thus frequency), its standing wave pattern will be spatially shifted relative to the first mode's pattern. A particularly relevant case is a mode whose antinodes fall where the first mode had its nodes. The nodes of $\cos^2\left(\frac{2\pi z}{\lambda}\right)$ are where $\frac{2\pi z}{\lambda} = \left(n + \frac{1}{2}\right)\pi$, so $z = \frac{\left(n+\frac{1}{2}\right)\lambda}{2} = \frac{(2n+1)\lambda}{4}$. The antinodes are at $z = \frac{n\lambda}{2}$. If another mode (mode 2) has its antinodes at the nodes of mode 1, it can utilize the population inversion that was left unsaturated by mode 1. The condition for this spatial complementarity is often simplified by saying a mode shifted by $\lambda/4$ (in terms of path length for its peaks relative to the original wave's peaks) can exist. More precisely, if the antinodes of a second potential lasing mode are located in the regions of minimal saturation (the nodes of the first mode), then this second mode can reach threshold by drawing on this "unused" inversion. This "implies multiple modes may co-exist even in homogeneous media." Recall that for a perfectly homogeneous medium *without* spatial hole burning, we expected strong mode competition leading to single-mode output. Spatial hole burning provides a mechanism that can undermine this, allowing several modes to lase simultaneously by drawing gain from different spatial regions of the active medium. Each mode "burns its own spatial holes" but leaves gain for other modes whose spatial patterns are different.

Page 58:

This page shows a diagram illustrating the "Spatial Hole Burning — Standing-Wave Saturation Pattern," specifically "(a) Standing-Wave Intensities."

The vertical axis is "Intensity $\frac{I(z)}{I_{\text{max}}}$" (Eye of zee divided by Eye max), ranging from 0.0 to 1.0.

The horizontal axis is "Position z," marked in units of wavelength λ: $\frac{1}{4}\lambda$, $\frac{1}{2}\lambda$, $\frac{3}{4}\lambda$, λ, $\frac{5}{4}\lambda$, etc., up to $\frac{9}{4}\lambda$.

Two thick gray vertical bars at the ends might represent the cavity mirrors, though their exact position relative to the λ markings isn't fully specified (e.g. if a mirror is at $z = 0$, we should have a node there). The plot extends a bit beyond 2λ.

Two sinusoidal intensity patterns are shown:

1. A solid blue curve, labeled "$I_1(z) \propto \cos^2\left(\frac{2\pi z}{\lambda}\right)$ (Saturating)". This represents the intensity profile of the primary lasing mode. It shows peaks (antinodes) at $z = 0$, $\frac{\lambda}{2}$, λ, $\frac{3\lambda}{2}$, 2λ, ... and nodes (zero intensity) at $z = \frac{\lambda}{4}$, $\frac{3\lambda}{4}$, $\frac{5\lambda}{4}$, $\frac{7\lambda}{4}$, ... (This assumes the cosine function starts at a maximum at $z = 0$. If $z = 0$ were a mirror, we'd expect a sine function for E, so $\sin^2$ for I, starting at zero). Let's assume the z-axis is relative to an antinode.

2. A dashed red curve, labeled "$I_2(z) \propto \sin^2\left(\frac{2\pi z}{\lambda}\right)$ (Exploiting)". This represents the intensity profile of a potential second mode that is spatially out of phase with the first one. Its peaks (antinodes) are precisely where the blue curve has its nodes (e.g., at $\frac{\lambda}{4}$, $\frac{3\lambda}{4}$, $\frac{5\lambda}{4}$, ...). Its nodes are where the blue curve has its antinodes.

The idea is that the blue "saturating" mode reduces the gain at its antinodes. The red "exploiting" mode can then lase by using the gain that remains at *its* antinodes (which were the nodes of the blue mode).

This figure clearly visualizes how two spatially distinct standing wave patterns can exist. If I_1 saturates the gain where it is strong, I_2 finds unsaturated gain where *it* is strong.

Page 59:

This is part (b) of the spatial hole burning illustration, showing the "(b) Population Inversion Profile $N(z)$" (Capital En of zee).

The axes are similar: vertical is "Population Inversion $N(z)$" (ranging from 0 up to N_0, the unsaturated inversion), and horizontal is "Position z" (in units of $\frac{\lambda}{4}$, $\frac{\lambda}{2}$, etc.).

We see:

* A dotted horizontal line at the top, labeled "N_0 (Initial Inversion)". This is the uniform population inversion before any lasing occurs.

A solid green curve, labeled "$N(z)$ (Saturated Inversion)". This shows the population inversion profile after* the first mode (I_1 from the previous slide, the $\cos^2$ wave) has started lasing and caused saturation. The green curve shows periodic dips. These dips, labeled "Spatial Hole," occur at the locations of the antinodes of I_1 (e.g., at $z = 0$, $\frac{\lambda}{2}$, λ, ...). At these positions, the strong intensity of I_1 has depleted the inversion. Between these dips, the green curve rises back up, approaching the original N_0 level. These peaks in the saturated inversion profile occur at the locations of the nodes of I_1 (e.g., at $z = \frac{\lambda}{4}$, $\frac{3\lambda}{4}$, ...). These are the "Unsaturated Regions."

This graph is the direct consequence of the intensity pattern $I_1(z)$. Where $I_1(z)$ is high, $N(z)$ is low (saturated). Where $I_1(z)$ is low, $N(z)$ remains high (unsaturated).

A second mode, like $I_2(z)$ whose antinodes align with these unsaturated regions (peaks in the green $N(z)$ curve), can then find sufficient gain to lase.

This clearly shows the spatial modulation of gain available to other modes due to the standing wave pattern of an existing mode.

Page 60:

Now we look at a "Quantitative Treatment of Spatial Hole-Burning Mode Spacing." This is Example 5.10 in some texts.

* "Suppose gain medium length L is short and is located a distance a (lowercase ay) from mirror M_1 (capital Em sub one)."

* "Condition for two standing waves of wavelengths λ_1 (lambda sub one) and λ_2 (lambda sub two) to have maxima separated by $\frac{\lambda}{p}$:" (lambda divided by p, where p is an integer).

This condition relates to how efficiently a second mode can utilize the gain left over by a first mode. If their antinodes are spatially shifted such that they sample different parts of the gain medium, they can coexist more easily. The $\frac{\lambda}{p}$ separation is a bit abstract here. A more direct condition is often related to the phase shifts within the gain medium.

The equation given is:

$$m\lambda_1 = a = \left(m + \frac{1}{p}\right)\lambda_2$$

Here, 'm' is an integer (a mode number related to how many half-wavelengths fit in distance a).

This equation seems to be setting up a condition where the gain medium at distance a experiences an antinode for λ_1 (if $a = \frac{m\lambda_1}{2}$ for field, or $a = m\lambda_1$

for some reference points). And for λ_2, the same location a corresponds to a slightly different phase $\left(m + \frac{1}{p}\right)$.

This is a somewhat specialized setup. The parameter 'p' (often $p = 2$ for $\frac{\lambda}{4}$ shift of intensity peaks, or $p = 4$ for fields) relates to the spatial relationship between modes that might allow them to coexist due to spatial hole burning. For $p = 2$, it means that at distance a, one mode has m half-wavelengths, and another has $m + \frac{1}{2}$ half-wavelengths, meaning one has an antinode and the other has a node (or vice-versa) within the gain medium if a is the location of the medium.

Let's take this as the starting point for relating frequencies.

Page 61:

Continuing from the condition $m\lambda_1 = a$ and $\left(m + \frac{1}{p}\right)\lambda_2 = a$:

* "where m is an integer."

* "Convert to frequency ($\nu = \frac{c}{\lambda}$):" From $m\lambda_1 = a$, we have $\lambda_1 = \frac{a}{m}$. So, $\nu_1 = \frac{c}{\lambda_1} = \frac{mc}{a}$.

From $\left(m + \frac{1}{p}\right)\lambda_2 = a$, we have $\lambda_2 = \frac{a}{\left(m+\frac{1}{p}\right)}$. So, $\nu_2 = \frac{c}{\lambda_2} = \frac{\left(m+\frac{1}{p}\right)c}{a}$.

These are the frequencies of two modes whose standing wave patterns have a specific spatial relationship (defined by 'p') within the gain medium located at 'a'.

* "Spacing between spatial-hole-burning modes:" This is $\delta\nu_{sp}$ (delta nu sub ess pee), the frequency difference $\nu_2 - \nu_1$:

$$\delta\nu_{sp} = \nu_2 - \nu_1 = \left[\frac{\left(m + \frac{1}{p}\right)c}{a}\right] - \left[\frac{mc}{a}\right]$$

$$\delta\nu_{sp} = \frac{c}{a}\left(m + \frac{1}{p} - m\right) = \frac{c}{a}\left(\frac{1}{p}\right)$$

So, $\delta\nu_{sp} = \frac{c}{ap}$.

This is a key result. It gives the frequency separation between two modes that are "spatially compatible" in the sense that one can use the gain left by the other due to spatial hole burning, for a gain medium at distance 'a' and a spatial shift factor 'p'.

* "Expressed relative to axial FSR $\delta\nu = \frac{c}{2\,d}$:" The axial FSR (Free Spectral Range) of the main cavity is $\delta\nu_{axial} = \frac{c}{2\,d}$. We want to relate $\delta\nu_{sp}$ to this.

Page 62:

The relationship between $\delta\nu_{sp}$ (delta nu sub ess pee) and the axial FSR $\delta\nu$ (delta nu, which is $\frac{c}{2\,d}$) is given by:

$$\delta\nu_{sp} = \frac{2\,d}{ap}\,\delta\nu$$

Let's verify this. We have $\delta\nu_{sp} = \frac{c}{ap}$ and $\delta\nu = \frac{c}{2\,d}$. So,

$$\frac{\delta\nu_{sp}}{\delta\nu} = \frac{\frac{c}{ap}}{\frac{c}{2\,d}} = \frac{c}{ap} \cdot \frac{2\,d}{c} = \frac{2\,d}{ap}.$$

Thus, $\delta\nu_{sp} = \frac{2\,d}{ap}\,\delta\nu$. This is correct.

This formula tells us that the characteristic frequency spacing for modes that can coexist due to spatial hole burning depends on the ratio of the total cavity length d to the distance a of the gain medium from a mirror, and on the factor p. If the gain medium is short and centrally located ($a \approx \frac{d}{2}$), and if $p = 2$ (corresponding to modes that are spatially orthogonal with respect to where they draw gain), then

$$\delta\nu_{sp} \approx \frac{2\,d}{\left(\frac{d}{2}\right) \cdot 2}\,\delta\nu = \frac{2\,d}{d}\,\delta\nu = 2\delta\nu.$$

This means the spatial hole burning modes would be spaced by twice the normal axial mode FSR. This is a common result for a short gain medium in the center of a cavity, allowing modes $q, q+2, q+4,$ ... to lase if the fundamental mode is q.

* Now, a condition for single-mode operation even with spatial hole burning: If the homogeneous gain width $\Delta\nu_{hom}$ is less than or equal to $\frac{2}{3}\left(\frac{d}{a}\right)\delta\nu$, only one spatial-hole-burning mode can survive. The term $\left(\frac{d}{a}\right)\delta\nu$ is not $\delta\nu_{sp}$. Let's re-examine

$$\delta\nu_{sp} = \frac{2\,d}{ap}\,\delta\nu.$$

If $p = 2$, $\delta\nu_{sp} = \frac{d}{a}\delta\nu$. The condition given is "If the homogeneous gain width $< \frac{2}{3}\delta\nu_{sp}$" (using $p = 2$ to get $\delta\nu_{sp} = \frac{d}{a}\delta\nu$), only one spatial-hole-burning mode can survive.

This implies that if the homogeneous linewidth (which determines the width of the *spectral* hole burned by a lasing mode) is narrower than about $\frac{2}{3}$ of the frequency spacing $\delta\nu_{sp}$ to the next "competing" spatial hole burning mode, then that next mode will fall too far into the spectrally burned hole of the first mode to lase. In other words, for a second mode to take advantage of the *spatial* hole, it must also find sufficient *spectral* gain. If the spectral hole burned by the first mode is wide enough to suppress gain at the frequency of the second potential spatial mode, then only one mode will lase. The factor $\frac{2}{3}$ is specific and might come from a detailed calculation of how much gain is needed. The general idea is: $\Delta\nu_{hom} < \delta\nu_{sp}$ is needed for robust single mode operation in this context.

Page 63:

Now we have "Example 5.10 — Numerical Illustration of Spatial Hole Burning." This corresponds to "Slide 19" by its internal numbering.

* "Given:" * "Cavity length $d = 100$ cm." * "Gain medium length $L = 0.1$ cm. This is indeed a short gain medium compared to d." * "Distance from mirror $a = 5$ cm. This is the distance of this short gain medium from one of the cavity mirrors." * "Desired separation index $p = 3$. This 'p' value determines the spatial relationship between the modes we are considering for coexistence due to spatial hole burning."

We need to calculate $\delta\nu$ (FSR) and $\delta\nu_{\text{sp}}$ (spatial hole burning mode spacing).

Page 64:

Let's perform the "Calculated quantities:"

1. First, the axial mode FSR, $\delta\nu$:

$$\delta\nu = \frac{c}{2\,d}$$

$$c = 3.00 \times 10^8 \text{ m/s} = 3.00 \times 10^{10} \text{ cm/s}$$

(since d is in cm).

$$d = 100 \text{ cm}$$

. So

$$2\,d = 200 \text{ cm}$$

.

$$\delta\nu = \frac{3.00 \times 10^{10} \text{ cm/s}}{200 \text{ cm}}$$

The slide shows:

$$\delta\nu = \frac{c}{2\,d} = \frac{3.00 \times 10^{8}}{2} \text{ Hz} = 150 \text{ MHz}$$

.

This calculation on the slide implicitly used $d = 1\text{m}$ not $d = 100 \text{ cm} = 1\text{m}$ in the denominator for

$$\frac{3 \times 10^{8}}{2 \times 1} = 1.5 \times 10^{8} \text{ Hz} = 150 \text{ MHz}$$

.

Let's re-do it with $d = 100 \text{ cm} = 1\text{m}$.

$$\delta\nu = \frac{3.00 \times 10^{8} \text{ m/s}}{2 \times 1.0 \text{ m}} = 1.50 \times 10^{8} \text{ Hz} = 150 \text{ MHz}$$

.

So, FSR $\delta\nu = 150$ MHz. This is correct.

2. Next, the spacing between spatial-hole-burning modes, $\delta\nu_{\text{sp}}$:

$$\delta\nu_{\text{sp}} = \frac{c}{ap}$$

.

$$a = 5 \text{ cm}$$

.

$$p = 3$$

.

So,

$$ap = 5 \text{ cm} \times 3 = 15 \text{ cm}$$

.

$$\delta\nu_{\text{sp}} = \frac{3.00 \times 10^{10} \text{ cm/s}}{15 \text{ cm}} = \frac{300}{15} \times 10^8 \text{ Hz} = 20 \times 10^8 \text{ Hz} = 2.0 \times 10^9 \text{ Hz}$$
$$= 2.0 \text{ GHz}$$

.

"delta nu sub ess pee equals $\frac{c}{ap}$ which is two hundred centimeters (this seems to be $2\,d$ written here instead of c) ... "

Let's look at the slide's formula:

$$\delta\nu_{\text{sp}} = \left(\frac{2\,d}{ap}\right)\delta\nu$$

.

This is the safer formula to use if $\delta\nu$ is already calculated.

$$2\,d = 200 \text{ cm}$$

.

$$a = 5 \text{ cm}$$

.

$$p = 3$$

.

$$\delta\nu = 150 \text{ MHz}$$

.

$$\delta\nu_{\text{sp}} = \frac{200 \text{ cm}}{(5 \text{ cm} \times 3)} \times 150 \text{ MHz}$$

.

$$\delta\nu_{\text{sp}} = \frac{200}{15} \times 150\ \text{MHz} = \frac{40}{3} \times 150\ \text{MHz} = 40 \times 50\ \text{MHz} = 2000\ \text{MHz}$$
$$= 2.0\ \text{GHz}$$

.

The slide shows:

$$\delta\nu_{\text{sp}} = \left(\frac{2\,d}{ap}\right)\delta\nu = \frac{200\ \text{cm}}{(5\ \text{cm} \times 3)} \times 150\ \text{MHz} \approx 2.0\ \text{GHz}$$

.

This calculation is correct. The spatial hole burning modes are separated by 2.0 GHz.

* "Interpretation:"

Page 65:

The interpretation of these calculated values is:

* "If the homogeneous width of the gain profile $\Delta\nu_{\text{hom}}$ (Delta nu sub hom) is less than 2.0 GHz, spatial hole burning will not allow more than one mode ⇒ single-mode operation feasible."

Let's analyze this. The characteristic spacing $\delta\nu_{\text{sp}}$ for modes that could coexist due to spatial hole burning is 2.0 GHz.

If a first mode is lasing, it burns both a spectral hole (of width roughly $\Delta\nu_{\text{hom}}$) and spatial holes.

For a second "spatial" mode to lase, it must find gain both spectrally and spatially.

If $\Delta\nu_{\text{hom}}$ is less than $\delta\nu_{\text{sp}}$ (2.0 GHz), it means that the spectral hole burned by the first lasing mode is narrower than the frequency gap to the next potential spatial hole burning mode.

This means that the next spatial mode (at a frequency $\nu_1 \pm \delta\nu_{\text{sp}}$) would fall outside (or on the very wings of) the spectral hole of the first mode. So it seems it *could* lase.

The statement on the slide is: "If $\Delta\nu_{\text{hom}} < 2.0$ GHz, spatial hole burning will *not* allow more than one mode." This implies that if $\Delta\nu_{\text{hom}}$ is narrow, the first mode is so dominant or efficient that it suppresses others.

This seems counterintuitive to the usual argument that if $\Delta\nu_{\text{hom}}$ is *small* compared to $\delta\nu_{\text{sp}}$, other modes can find gain.

Perhaps the logic is: for a mode at $\nu_1 + \delta\nu_{\text{sp}}$ to lase, it needs gain. The mode at ν_1 will burn a spectral hole of width $\sim \Delta\nu_{\text{hom}}$ around ν_1. If $\Delta\nu_{\text{hom}} < \delta\nu_{\text{sp}}$, then the mode at $\nu_1 + \delta\nu_{\text{sp}}$ is spectrally distinct.

The condition for the *second* mode (at $\nu_1 + \delta\nu_{\text{sp}}$) to be suppressed by the *first* mode (at ν_1) would normally be if the *spectral* hole burned by ν_1 is wide enough to reduce the gain at $\nu_1 + \delta\nu_{\text{sp}}$ below threshold. This happens if $\Delta\nu_{\text{hom}}$ is comparable to or larger than $\delta\nu_{\text{sp}}$.

So, if $\Delta\nu_{\text{hom}} > \delta\nu_{\text{sp}}$, the spectral hole is broad, and the first mode would suppress the second one, leading to single mode operation.

If $\Delta\nu_{\text{hom}} < \delta\nu_{\text{sp}}$, the spectral hole is narrow, and the second mode at $\nu_1 + \delta\nu_{\text{sp}}$ might be able to lase because it's spectrally far enough away from the hole, AND it benefits from the spatial hole. This would lead to *multimode* operation.

The slide's statement: "If $\Delta\nu_{\text{hom}} < 2.0$ GHz, spatial hole burning will not allow more than one mode $\Rightarrow$ single-mode operation feasible."

This means if the homogeneous width is narrower than the spatial mode separation, we get single mode. This is the opposite of my reasoning above.

Let's re-think. The condition for a second mode to lase is that the gain available to it is above threshold. Spatial hole burning *provides* gain spatially. Spectral hole burning *removes* gain spectrally.

If $\Delta\nu_{\text{hom}}$ is very small, the spectral hole is very narrow. A mode at $\nu_1 + \delta\nu_{\text{sp}}$ is far from this narrow spectral hole, so spectrally it sees lots of gain. Spatially, it also sees gain. So it should lase. This leads to multimode.

If $\Delta\nu_{\text{hom}}$ is large (e.g., $\Delta\nu_{\text{hom}} \approx \delta\nu_{\text{sp}}$), the spectral hole is wide. The mode at $\nu_1 + \delta\nu_{\text{sp}}$ now falls within this wide spectral hole. So it is suppressed. This leads to single mode.

Therefore, the statement on the slide seems to be reversed. It should likely be: "If the homogeneous width $\Delta\nu_{\text{hom}}$ is *greater* than ~ 2.0 GHz (i.e., comparable to or larger than $\delta\nu_{\text{sp}}$), then the spectral hole burned by the first mode is wide enough to suppress other spatial hole burning modes, potentially leading to single-mode operation." Or, "If $\Delta\nu_{\text{hom}} < 2.0$ GHz, multiple modes can coexist due to spatial hole burning."

Let's assume there might be a nuance I'm missing or a specific context from the textbook. However, standard understanding suggests narrower $\Delta\nu_{\text{hom}}$ (relative to mode spacing) facilitates multimode behavior when spatial hole burning is present. The slide's phrasing might imply that if $\Delta\nu_{\text{hom}}$ is very small, the system is so "homogeneous" spectrally that one mode dominates completely, despite spatial effects. This would only be true if spatial hole burning was not effective enough to allow the second mode.

Let's take the slide's statement at face value for now: if $\Delta\nu_{\text{hom}} < 2.0$ GHz (our calculated $\delta\nu_{\text{sp}}$), then single-mode operation is feasible due to spatial hole burning not allowing more than one mode. This would imply that the primary mode saturates the spatially available gain so effectively that even if other modes are spectrally clear, they can't find enough *spatial* gain. This would be a strong effect of the primary mode.

Page 66:

Now we discuss methods for "Eliminating Spatial Hole Burning — Ring Lasers." Spatial hole burning is a consequence of the standing wave pattern in a linear cavity. If we can eliminate the standing wave, we can eliminate spatial hole burning.

"In a unidirectional ring cavity the field is purely travelling wave."

A ring laser uses three or more mirrors to guide the light beam in a closed loop. If the laser is made to operate unidirectionally (i.e., light travels only clockwise, or only counter-clockwise, but not both), then the field inside the cavity is a "travelling wave," not a standing wave. Optical isolators or other non-reciprocal elements are often used to enforce unidirectional operation.

"Intensity I is uniform along the propagation axis ⇒ no nodes and antinodes."

For a pure travelling wave, the intensity is ideally constant along the path of the beam within the gain medium (ignoring absorption or gain effects for a moment, just focusing on the wave pattern). There are no fixed positions of zero intensity (nodes) or maximum intensity (antinodes) like in a standing wave.

"Result:"

- "Entire gain medium experiences the same saturation ⇒ no spatial holes." Since the intensity is uniform spatially (along the beam path within the gain medium), every part of the gain medium experiences the same light intensity. Therefore, gain saturation occurs uniformly throughout the volume of the gain medium that interacts with the beam. There are no unsaturated regions left behind due to a standing wave pattern.

- "Higher overall extraction efficiency; favoured in high-power lasers." Because the entire gain medium contributes to the amplification (it's all being saturated and contributing photons to the lasing mode), ring lasers can often achieve higher extraction efficiency – meaning they can convert more of the stored energy in the gain medium into laser output. This is particularly advantageous for high-power laser systems where maximizing output is crucial. By avoiding "dead zones" of unused inversion at the nodes, more of the pumped volume can contribute.

Page 67:

Let's look at "Practical ring implementations:"

* "Four-mirror folded ring (bow-tie)." This is a common configuration for ring lasers. Four mirrors are arranged to form a path that looks like a bow-tie. This geometry can be designed to include a focus within the gain medium and can accommodate other intracavity elements.

* "Fibre ring lasers." Optical fibers can also be used to create ring cavities. A loop of optical fiber, with a section of doped fiber acting as the gain medium, and a fiber coupler to extract output, can form a very compact and stable ring laser. These are widely used in telecommunications and sensing.

* "Unidirection enforced with optical isolators or Faraday rotators." To ensure that the ring laser operates with a travelling wave in only one direction (and not as a standing wave cavity due to reflections allowing both directions), a non-reciprocal element is usually inserted into the cavity. An "optical isolator" allows light to pass in one direction but blocks or highly attenuates it in the reverse direction. This is often based on the "Faraday effect," where a magnetic field applied to a special material (like a Faraday rotator made of TGG crystal or YIG) rotates the plane of polarization of light. Combined with polarizers, this can create one-way transmission. By forcing unidirectional operation, standing waves are suppressed, and thus spatial hole burning is eliminated. This then allows the mode competition in

a homogeneously broadened medium to effectively lead to single-mode operation if desired (or at least, it removes one mechanism for multimode operation).

Page 68:

Now we summarize "Homogeneous vs Inhomogeneous — Consequences for Multimode Operation."

"Homogeneous gain:"

"Strong mode competition through shared population."

As we discussed, in a homogeneously broadened medium, all atoms are identical and share the same gain profile. If one mode starts lasing and saturates the gain, it affects the gain available to all other potential modes. This leads to strong competition.

"Tends towards single-mode output *if* spatial hole burning and technical noise can be suppressed."

The asterisk is important. The natural tendency of a purely homogeneous system is single-mode output. However, if spatial hole burning is present (as in a linear cavity), it can enable multimode operation by allowing different modes to access different spatial regions of gain. "Technical noise" (like vibrations causing slight changes in cavity length, or pump fluctuations) can also sometimes cause the lasing mode to hop between different cavity modes, or allow multiple modes to lase if the competition isn't perfectly decisive. If these factors are controlled (e.g., by using a ring cavity to eliminate spatial hole burning, and by stabilizing the laser), then homogeneous gain media are good candidates for single-frequency operation.

"Inhomogeneous gain:"

Page 69:

Continuing with the consequences for inhomogeneous gain:

* "Minimal competition; each mode taps a distinct sub-ensemble." In an inhomogeneously broadened medium, different cavity modes can interact with different groups (sub-ensembles) of atoms/molecules, those whose individual Doppler-shifted or site-shifted resonant frequencies align with the specific cavity mode. Since these sub-ensembles are largely independent (a mode at ν_1 saturates atoms around ν_1, leaving atoms around ν_2 for a different mode), the competition between modes is much weaker than in the homogeneous case. One mode burning a spectral hole at its frequency doesn't necessarily prevent another mode, far away in frequency, from lasing using a different group of atoms.

* "Many axial and transverse modes can oscillate simultaneously." As a result of this weaker competition and the availability of distinct gain packets, inhomogeneously broadened lasers (like He-Ne or Argon ion lasers) often operate on multiple axial modes simultaneously, especially if no mode selection elements are used. They can also support multiple transverse modes (TEM_{mn} modes) if the cavity design allows.

* "Real media are mixed; the ratio $\delta\nu/\Delta\nu_{\text{hom}}$ (delta nu divided by Delta nu sub hom) dictates competition strength." In reality, no gain medium is perfectly homogeneous or perfectly inhomogeneous. There's often a mix. The homogeneous linewidth $\Delta\nu_{\text{hom}}$ of the individual atomic packets is finite. The cavity axial modes are spaced by $\delta\nu$ (the FSR).

The ratio of these two, $\delta\nu/\Delta\nu_{\text{hom}}$, is crucial.

If $\Delta\nu_{\text{hom}} \gg \delta\nu$ (homogeneous linewidth is much larger than mode spacing), then several cavity modes fall within one homogeneous packet. These modes will compete strongly, as in a purely homogeneous system. The medium behaves more homogeneously.

If $\Delta\nu_{\text{hom}} \ll \delta\nu$ (homogeneous linewidth is much smaller than mode spacing), then each cavity mode essentially interacts with its own distinct group of atoms. The modes are well separated compared to the spectral hole width. Competition is weak. The medium behaves more inhomogeneously.

So, this ratio determines whether the laser behaves more like a homogeneous or inhomogeneous system in terms of mode competition.

* "Transverse modes ($\mathrm{TEM}_{m,n}$) complicate the spectrum because they have different longitudinal frequencies in non-confocal resonators." We've mostly focused on axial modes (q). Transverse ElectroMagnetic modes, TEM_{mn}, describe the intensity distribution in the plane perpendicular to the cavity axis. The fundamental mode is TEM_{00} (a Gaussian beam). Higher-order modes (TEM_{01}, TEM_{10}, etc.) have more complex spatial patterns.

In general, for a non-confocal resonator (most resonators are not perfectly confocal), the resonant frequencies of these transverse modes are not degenerate with the axial modes. The frequency of a TEM_{mnq} mode is given by

$$\nu_{mnq} = \frac{c}{2\,d}\left[q + \frac{1}{\pi}(m + n + 1)\arccos\sqrt{g_1\,g_2}\right]$$

where g_1 and g_2 are cavity stability parameters.

This means that for a given axial mode number q, different transverse modes (different m, n) will have slightly different resonant frequencies. This adds further complexity to the laser output spectrum, effectively creating more "lines" unless measures are taken to suppress these higher-order transverse modes.

Page 70:

This brings us to "Designing Cavities to Suppress Higher-Order Transverse Modes." For many applications, especially in spectroscopy or when a high-quality beam profile is needed, operation in the fundamental TEM_{00} mode is desired. Higher-order transverse modes have more complex spatial profiles and typically larger divergence.

* "Key tools:" How can we achieve TEM_{00} operation?

1. "Choose mirror radii such that the resonator is stable only for TEM_{00}." The stability conditions for a resonator depend on the mirror curvatures and separation. It's possible to design a cavity that is on the edge of a stability zone, where it might be stable for the lowest-loss TEM_{00} mode but unstable (and thus highly lossy) for higher-order modes, which tend to be spatially larger. This is a more advanced technique.

2. "Insert intracavity apertures matched to the Gaussian beam waist." This is the most common method. The TEM_{00} mode has the smallest spot size within the cavity. Higher-order transverse modes are spatially larger. By inserting an aperture (a pinhole or adjustable iris) inside the cavity, typically at a location where the TEM_{00} beam waist is, one can introduce significant losses for the larger higher-order modes, effectively preventing them from lasing. The aperture is sized to allow the TEM_{00} mode to pass with minimal loss, while clipping and attenuating the higher-order modes.

3. "Employ unstable or confocal resonator geometries when high power is needed but single transverse mode is tolerated." This point seems a bit contradictory. "Single transverse mode is tolerated" sounds like higher orders are okay. Perhaps it means "when high power is needed AND single transverse mode operation is desired." **Unstable resonators** are often used for high-power lasers with large gain volumes. They don't confine the beam in the traditional sense but allow it to expand and fill the gain volume, with output typically taken via diffraction around one of the mirrors. They can produce good beam quality (though not necessarily pure TEM_{00} in the same sense as stable resonators) and high power. **Confocal resonators** ($d = R_1 = R_2 = R$) have the property that the resonant frequencies of all transverse modes are degenerate (or nearly so for a perfectly aligned

confocal cavity). This means all TEM$_{mnq}$ modes for a given $q + (m + n)$ would have the same frequency. This can lead to a cleaner spectrum in some senses but doesn't inherently suppress higher-order modes; they just overlap in frequency with the TEM$_{00}$ modes of different q. The point about "single transverse mode is tolerated" might be a typo for "is desired" or "is critical". If single transverse mode is merely "tolerated", then one might not need to do much. If it's "desired," apertures are key for stable resonators. Unstable resonators are a different class often used for high gain systems to extract power efficiently with good beam quality, often resembling a single transverse mode after propagation.

Page 71:

The benefits and implications of achieving single transverse mode operation (typically TEM$_{00}$):

"Benefit: a Gaussian beam with minimal divergence, higher spatial coherence."

The TEM$_{00}$ mode has a Gaussian intensity profile. A pure Gaussian beam has the property of minimal beam divergence for a given beam waist size; it's diffraction-limited. This means it can be focused to the smallest spot and will spread out the least as it propagates.

It also exhibits high "spatial coherence." This means that the phase of the wavefront is well-defined and correlated across the beam's cross-section. High spatial coherence is essential for applications like interferometry, holography, and anything requiring a clean, predictable wavefront.

"Spectral implication: still potentially multimode in frequency unless axial mode selection is also enforced."

This is a crucial reminder. Suppressing higher-order *transverse* modes (making the laser operate in TEM$_{00}$) only cleans up the spatial profile of the beam.

It does *not* automatically mean the laser will operate on a single *axial* mode (i.e., a single frequency).

A TEM_{00} laser can still have multiple axial modes $(q, q+1, q+2, \ldots)$ lasing simultaneously if they all fall under the net gain curve.

So, if true single-frequency operation is required, one must employ techniques to select a single axial mode (like intracavity etalons, as discussed for dye lasers) *in addition* to ensuring TEM_{00} operation. The two are separate (though related) aspects of controlling the laser output.

Page 72:

Let's look at "Example 5.11 — HeNe Laser Detailed Mode Count." This is Slide 23.

We'll revisit the Helium-Neon laser example with some specific parameters to count how many modes might actually lase.

* "Parameters recap:" * " $\lambda = 632.8\,\text{nm}$." (lambda equals six hundred thirty-two point eight nanometers). This is the classic red He-Ne laser wavelength. * "Doppler width $1.5\,\text{GHz}$ (FWHM)." This is the Full Width at Half Maximum of the inhomogeneous Doppler-broadened gain profile of the neon transition. Let's call this $\Delta\nu_{\text{D}}$. * "Cavity length $d = 1.0\,\text{m}$." (dee equals one point zero meters). * "Axial mode spacing $\delta\nu = 150\,\text{MHz}$." (delta nu equals one hundred fifty megahertz). This is the FSR for a 1m cavity:

$$\delta\nu = \frac{c}{2\,d} = \frac{3 \times 10^8\ \text{m/s}}{2 \times 1\ \text{m}} = 1.5 \times 10^8\ \text{Hz} = 150\ \text{MHz}$$

This is consistent.

Page 73:

Continuing with the HeNe laser example:

* "Gain above threshold $\approx 1.2\,\text{GHz}$ wide." This is the width of the frequency region where the small-signal gain curve is above the loss line (like the

shaded region in the graph on page 28 or 35). Let's call this $\Delta\nu_{\text{lase}}$. It's given as $1.2\,\text{GHz}$. This is narrower than the full Doppler width of $1.5\,\text{GHz}$, which is expected, as only the central part of the gain curve will typically be above threshold.

* "Potential axial modes $N \approx 8$." We can calculate this as $N = \frac{\Delta\nu_{\text{lase}}}{\delta\nu} = \frac{1.2\,\text{GHz}}{150\,\text{MHz}} = \frac{1200\,\text{MHz}}{150\,\text{MHz}} = \frac{120}{15} = 8$. So, approximately 8 axial cavity modes fall within the region where gain exceeds loss.

Now, we need to consider the "Homogeneous linewidth contributions:"

Even though the He-Ne laser line is primarily Doppler (inhomogeneously) broadened, each individual neon atom (or velocity group of atoms) has an underlying homogeneous linewidth, $\Delta\nu_{\text{hom}}$. This $\Delta\nu_{\text{hom}}$ determines the width of the spectral hole that a lasing mode will burn.

Several factors contribute to $\Delta\nu_{\text{hom}}$:

* "Natural width $20\,\text{MHz}$." This arises from the finite radiative lifetime of the upper laser level (related to the Einstein A_{21} coefficient). $\Delta\nu_{\text{natural}} = \frac{A_{21}}{2\pi}$. A typical value is given as $20\,\text{MHz}$.

* "Pressure width $\approx 20\,\text{MHz}$." This is due to collisional broadening. Collisions between neon atoms (and with helium atoms) interrupt the phase of the emission, leading to broadening. At typical He-Ne laser pressures, this contributes about $20\,\text{MHz}$.

* "Power broadening (example $I/I_{\text{s}} = 10$): increases to $\approx 100\,\text{MHz}$." If a mode is lasing with significant intensity, it will power-broaden its own homogeneous linewidth. The power-broadened homogeneous linewidth $\Delta\nu_{\text{hom}}'$ is approximately $\Delta\nu_{\text{hom}}' \approx \Delta\nu_{\text{hom}} \times \sqrt{1 + I/I_{\text{s}}}$. The "bare" $\Delta\nu_{\text{hom}}$ (from natural + pressure broadening) would be around $20\,\text{MHz} + 20\,\text{MHz} = 40\,\text{MHz}$. If $I/I_{\text{s}} = 10$ (strong saturation), then $\sqrt{1 + 10} = \sqrt{11} \approx 3.317$. So, $\Delta\nu_{\text{hom}}' \approx 40\,\text{MHz} \times 3.317 \approx 132.7\,\text{MHz}$. The slide says it "increases to $\approx$

$100\,\mathrm{MHz}$." This is in the same ballpark and suggests the example uses specific values for I_s and I, or a slightly different base $\Delta\nu_{hom}$. Let's use $\Delta\nu_{hom}' \approx 100\,\mathrm{MHz}$ as the effective homogeneous linewidth of a lasing mode under these conditions.

Page 74:

Now, we compare the effective homogeneous linewidth $\Delta\nu_{hom}'$ with the axial mode spacing $\delta\nu$.

- We found $\Delta\nu_{hom}'$ (the power-broadened homogeneous width, i.e., the width of the spectral hole burned by a lasing mode) to be around 100 MHz.
- We found the axial mode spacing $\delta\nu$ to be 150 MHz.
- "Since homogeneous width $< \delta\nu$ ($\Delta\nu_{hom}' < \delta\nu$, i.e., 100 MHz $<$ 150 MHz), modes do not overlap strongly ⇒ simultaneous oscillation of several independent modes yields stable multi-line spectrum (Fig. 5.27a)."

This is the key comparison for an inhomogeneously broadened laser like He-Ne. The spectral hole burned by one lasing mode (width ~ 100 MHz) is narrower than the spacing to the next cavity mode (150 MHz away).

This means that if mode q is lasing, it burns a hole around its frequency ν_q. The adjacent modes ν_{q-1} and ν_{q+1} are 150 MHz away. They fall largely *outside* this 100 MHz wide spectral hole.

Therefore, these adjacent modes can still find unsaturated (or less saturated) atoms in the inhomogeneous Doppler profile to provide them with gain.

Because the spectral holes are relatively isolated and don't strongly overlap and suppress neighbors, "simultaneous oscillation of several independent modes" is possible.

The laser will likely operate on several of the ~ 8 potential axial modes. Each mode will burn its own ~ 100 MHz wide spectral hole in the ~ 1.2 GHz wide gain-above-threshold profile.

This results in a "stable multi-line spectrum." The output will consist of several discrete frequencies, separated by 150 MHz. "Fig. 5.27a" in the textbook would presumably show such a spectrum.

This is characteristic behavior for a He-Ne laser: it's often multimode unless specific measures are taken to force single-mode operation (which is harder for inhomogeneously broadened lasers).

Page 75:

Let's consider another example:

Slide 24: Example 5.12 — Argon-Ion Laser Multimode Dynamics

Argon-ion lasers are another important class of gas lasers, capable of higher powers than He-Ne, often operating on blue or green lines.

"Key numerical data:"

"Doppler width $8–10$ GHz." This is the inhomogeneous width of the gain lines. This is significantly broader than 1.5 GHz for He-Ne.

"Resonator length $d = 1.2$ m $\Rightarrow \delta\nu = 125$ MHz."

Let's check the FSR:

$$\delta\nu = \frac{c}{2\,d} = \frac{3 \times 10^8 \text{ m/s}}{2 \times 1.2 \text{ m}} = \frac{3 \times 10^8}{2.4} \text{ Hz} = 1.25 \times 10^8 \text{ Hz} = 125 \text{ MHz}$$

This is correct. So, the cavity modes are spaced 125 MHz apart.

"Homogeneous width $\gg \delta\nu$ due to:" ($\Delta\nu_{\text{hom}}$ is much greater than $\delta\nu$).

This is a crucial difference from the He-Ne case we just saw (where $\Delta\nu_{\text{hom}}'$ was $< \delta\nu$). Why is $\Delta\nu_{\text{hom}}$ so large in an Argon-ion laser?

"Electron-ion pressure broadening."

Argon-ion lasers operate with a high-current electrical discharge, creating a hot, dense plasma. In this plasma, collisions between the argon ions and electrons are very frequent and energetic. These collisions cause significant broadening of the homogeneous linewidth of the lasing transitions. This "pressure broadening" (or more accurately, Stark broadening from the microfields of charged particles and direct collisional dephasing) can make $\Delta\nu_{\text{hom}}$ very large, often several hundred MHz or even into the GHz range, depending on discharge conditions.

Page 76:

Continuing with the Argon-ion laser:

* "High intracavity power (10–100 W) ⇒ strong power broadening." Argon-ion lasers can generate high continuous-wave powers. The intensity inside the laser cavity can be extremely high (kilowatts or megawatts per cm^2). This leads to a very large saturation parameter I/I_s, and thus very strong power broadening of the already large homogeneous linewidth. So, the effective $\Delta\nu_{\text{hom}}'$ can become extremely broad, potentially many GHz.

* "Resulting behaviour:"

* " $\Delta\nu_{\text{hom}} \gg \delta\nu \Rightarrow$ intense mode competition." (Here, $\Delta\nu_{\text{hom}}$ should be understood as $\Delta\nu_{\text{hom}}'$, the power-broadened homogeneous width). If the effective homogeneous width $\Delta\nu_{\text{hom}}'$ is much larger than the axial mode spacing $\delta\nu$ ($125\,\text{MHz}$ in this example), it means that the spectral hole burned by one lasing mode is broad enough to encompass and affect many adjacent cavity modes. For example, if $\Delta\nu_{\text{hom}}'$ is $1\,\text{GHz}$, it covers $\frac{1000}{125} = 8$ cavity modes. This situation leads to "intense mode competition" because many cavity modes are essentially trying to draw gain from the same homogeneously broadened packet of ions. The laser starts to behave more like a homogeneously broadened system, where one mode (or a few closely-knit modes) might try to dominate.

* "However, technical perturbations (cavity length jitter, discharge noise) constantly redistribute gain ⇒ observed spectrum fluctuates randomly (Fig. 5.27b)." While the large $\Delta\nu_{\text{hom}}'$ might suggest a tendency towards fewer modes due to competition, Argon-ion lasers are notorious for having unstable, fluctuating multimode spectra. The "technical perturbations" are significant:

* **Cavity length jitter:** Vibrations, thermal drifts cause the cavity length d to change slightly, which shifts all the cavity mode frequencies $\nu_{\text{q}} = \frac{qc}{2\,d}$.

* **Discharge noise:** The plasma discharge is inherently noisy, leading to fluctuations in temperature, ion density, and thus gain.

These perturbations "constantly redistribute gain" among the competing modes. No single mode can stably dominate. The laser might jump between different sets of modes, or the relative intensities of the lasing modes might fluctuate rapidly. The "observed spectrum fluctuates randomly." If you look at the output of such a laser with a fast spectrometer, you would see the pattern of lasing modes changing over time. "Fig. 5.27b" in the textbook would likely show this kind of unstable, "spiky" spectrum, perhaps averaged over time to look like a broad envelope but instantaneously quite different. This makes free-running Argon-ion lasers less suitable for applications requiring high spectral stability or single-frequency operation without active stabilization measures (like an intracavity etalon, which is often used).

Page 77:

Let's consider a third type of laser: "Slide 25: Example 5.13 — Dye Laser Spectral Features."

Dye lasers use organic dye molecules in a liquid solvent as the gain medium. They are known for their broad tunability.

• "Liquid dye medium: gain FWHM of order 2×10^{13} Hz (≈ 20 nm at 600 nm)." "gain Full Width at Half Maximum of order two times ten to the thirteenth Hertz." This is 20 THz. This is an extremely broad gain bandwidth, similar to what we saw for "solid-state/dye laser broad gain" earlier (page 38 had 1×10^{13} Hz). A bandwidth of 20 THz at a center wavelength of 600 nm corresponds to a wavelength range of about 20 nm.

$$\Delta\lambda \approx \frac{\lambda^2 \Delta\nu}{c} = \frac{(600\,\text{nm})^2 \cdot (2 \times 10^{13}\,\text{Hz})}{3 \times 10^8\,\text{m/s}} = \frac{(3.6 \times 10^{-13}\,\text{m}^{2)}) \cdot (2 \times 10^{13}\,\text{Hz})}{3 \times 10^8\,\text{m/s}}$$
$$= \frac{7.2\,\text{m}}{3 \times 10^8\,\text{m/s}} = 2.4 \times 10^{-8}\,\text{m} = 24\,\text{nm}$$

So, ≈ 20 nm is correct.

• "With d = 0.75 m, axial modes N ~ 10^5." (d is cavity length, N is number of modes) FSR

$$\delta\nu = \frac{c}{2\,d} = \frac{3 \times 10^8\,\text{m/s}}{2 \times 0.75\,\text{m}} = \frac{3 \times 10^8}{1.5}\,\text{Hz} = 2 \times 10^8\,\text{Hz} = 200\,\text{MHz}$$

Number of modes

$$N = \frac{\Delta\nu_{\text{gain}}}{\delta\nu} = \frac{2 \times 10^{13}\,\text{Hz}}{2 \times 10^{8}\,\text{Hz}} = 10^{13-8} = 10^{5}$$

So, about one hundred thousand axial modes fit under this gain curve!

- "Homogeneous profile dominates, yet refractive index fluctuations in the liquid cause mode coupling ⇒ time-dependent multimode emission." Dye laser transitions in liquid solutions are primarily homogeneously broadened. The rapid interactions (collisions, solvent relaxation) around the dye molecules ensure that all molecules share energy quickly and have a similar response. The homogeneous linewidth $\Delta\nu_{\text{hom}}$ can be quite large, on the order of THz itself (though usually smaller than the full gain bandwidth). Because $\Delta\nu_{\text{hom}}$ is typically much larger than $\delta\nu$ (THz vs 200 MHz), we would expect strong mode competition, potentially leading to few modes or single-mode operation if spatial hole burning is managed (e.g., in a ring dye laser, or with a jet stream for the dye). However, a complication arises: "refractive index fluctuations in the liquid cause mode coupling." The liquid dye solution is subject to turbulence, thermal gradients, and micro-bubbles, especially in the pumped region (often a high-intensity focused laser beam is used for pumping). These effects cause rapid, random fluctuations in the local refractive index of the dye solution. These refractive index fluctuations act like a time-varying phase perturbation inside the cavity. They can scatter light between different cavity modes ("mode coupling") and effectively disrupt the clean mode competition that would otherwise occur. This leads to "time-dependent multimode emission." Even if the laser tries to settle on one mode, these fluctuations can kick it into other modes or allow several modes to lase simultaneously in an unstable manner.

- "Pulsed dye lasers:" Many dye lasers are pumped by pulsed lasers (e.g., Nd:YAG, excimer lasers) and therefore operate in a pulsed mode.

Page 78:

Continuing with pulsed dye lasers:

* "Gain builds up in μs-scale (microsecond scale); spectrum integrates over hopping—final bandwidth ≈ 1 nm."

In a pulsed dye laser, the pump pulse might last nanoseconds to microseconds. During this time, the gain builds up, lasing starts, and the laser might "hop" between different sets of modes due to the refractive index fluctuations and the rapid evolution of conditions.

If we measure the output spectrum integrated over the entire laser pulse, we don't see the instantaneous sharp modes, but rather a broadened envelope. This time-integrated "final bandwidth" can be around 1 nm.

A 1 nm bandwidth at 600 nm corresponds to a frequency bandwidth of:

$$\Delta\nu = \frac{c\,\Delta\lambda}{\lambda^2} = \frac{(3\times 10^8\ \text{m/s})\,(1\times 10^{-9}\ \text{m})}{(600\times 10^{-9}\ \text{m})^2} = \frac{3\times 10^{-1}}{3.6\times 10^{-13}}\ \text{Hz} \approx 0.83\times 10^{12}\ \text{Hz} \approx 0.8\ \text{THz}.$$

So, even though the gain bandwidth is ~ 20 THz, and there are 10^5 modes, the time-averaged output of a simple pulsed dye laser might be around 0.8 THz wide (1 nm). This is still very broad for spectroscopy.

* "To obtain tunable single-mode operation:"

Given this natural tendency for broad, multimode emission, if we need narrow, tunable, single-frequency output from a dye laser (which is one of their key applications due to broad gain), we must implement aggressive mode selection techniques.

* "Introduce diffraction grating, prism, or birefringent filter inside cavity (Sect. 5.4)."

These are intracavity tuning elements that provide wavelength-dependent loss, effectively narrowing the gain bandwidth to select a much smaller region of frequencies.

* **Diffraction grating:** Often used in Littrow or Littman-Metcalf configuration, it acts like a frequency-selective mirror. * **Prism:** Can be used for tuning, though often provides broader selection. Multiple prisms (Brewster-angle prisms) can be used for dispersion without high loss. * **Birefringent filter (Lyot filter):** This consists of one or more birefringent plates between polarizers. It acts as a wavelength-dependent transmission filter, with multiple transmission peaks. By rotating the plates, the passband can be tuned.

These elements provide coarse to medium tuning. For true single-mode operation, one typically also needs one or more intracavity etalons, as mentioned in Section 5.4 (which we'll cover later, or refers to a previous/next section in the text). The combination of these elements restricts lasing to a single axial and transverse mode.

Page 79:

This page shows a graph illustrating "Pulsed Dye Laser Spectral Features."

The horizontal axis is "Wavelength (nm)," ranging from $580\,\mathrm{nm}$ to $620\,\mathrm{nm}$, centered around $600\,\mathrm{nm}$.

There are several features depicted:

1. A very broad, light blue, bell-shaped curve representing the "Gain Profile FWHM $\approx 20\,\mathrm{nm}$." This is the overall gain bandwidth of the dye, spanning roughly from $590\,\mathrm{nm}$ to $610\,\mathrm{nm}$ at half maximum. This corresponds to the $20\,\mathrm{THz}$ frequency bandwidth.

2. Inside this broad gain profile, there is a much narrower, black, spiky structure. This is labeled "Integrated Lasing Spectrum FWHM $\approx 1\,\mathrm{nm}$." This represents the time-averaged output of the pulsed dye laser if only coarse tuning elements (or none) are used. It's much narrower than the gain profile, but still broad ($1\,\mathrm{nm}$, which we found is $\sim 0.8\,\mathrm{THz}$). This black region is not smooth but shown with many fine spikes and gaps.

3. Annotations within this black region point to "Clusters of modes" and "Gaps." The "Clusters of modes" suggests that the output isn't a single continuous band but is composed of groups of cavity modes that happen to lase during the pulse. The "Gaps" indicate regions where, even within this $1\,\mathrm{nm}$ envelope, there might be no laser output.

4. A note at the bottom explains: "Clusters/gaps due to mode coupling, transient effects, and/or unintended etalon effects." * **Mode coupling:** Caused by refractive index fluctuations in the liquid dye, as discussed. * **Transient effects:** In pulsed operation, the populations and intensity are rapidly changing, so steady-state conditions may not be reached. This can lead to complex spectral dynamics. * **Unintended etalon effects:** Sometimes, other parallel surfaces within the laser cavity (like the windows of the dye cell, or even a misaligned etalon) can act as weak etalons, creating unwanted modulation in the spectrum, leading to clusters and gaps.

This graph vividly illustrates the challenge: how to go from the broad $20\,\mathrm{nm}$ gain profile, through an uncontrolled $\sim 1\,\mathrm{nm}$ multimode output, to a desired single, narrow frequency for spectroscopy. This requires the sophisticated intracavity elements mentioned.

Page 80

We now address "Non-Uniform Intensity Distribution Within a Multimode Band." This is Slide 26.

Even when a laser is operating multimode over a certain spectral band, the intensity is not necessarily smooth across that band.

"Laser spectral density $I_{\mathrm{L}}(\nu, t)$ is not smooth — contains peaks and holes:"

If multiple axial modes are lasing simultaneously, their phases and amplitudes can fluctuate, or they can interfere.

The instantaneous spectral density (intensity as a function of frequency and time) is given as a sum over modes k:

$$I_L(\omega, t) = \sum_k A_k(t)\cos^2[\omega_k t + \Phi_k(t)]$$

Let's clarify this formula. Usually, the intensity of mode k is $A_k(t)$. The $\cos^2[\omega_k t + \Phi_k(t)]$ term seems to imply some kind of interference beating or modulation, but if $A_k(t)$ is the intensity envelope of mode k at frequency ω_k, then the spectral density would be a sum of narrow peaks centered at each ω_k.

$$I_L(\omega, t) = \sum_k A_k(t)L(\omega - \omega_k, t)$$

where L is the lineshape of mode k.

The formula $A_k(t)\cos^2[\omega_k t + \Phi_k(t)]$ looks more like the electric field or a component of it, squared.

If ω_k are the mode frequencies and $\Phi_k(t)$ are their phases, then if these modes interfere on a detector, they can produce beat frequencies.

However, if this $I_L(\omega, t)$ is meant to be the *spectral* intensity, it would be a series of peaks at ω_k. The $\cos^2$ term is unusual for spectral density itself.

Perhaps it refers to the fact that even within what appears to be a continuous band, if you resolve it finely, it's made of discrete modes, and their relative strengths $A_k(t)$ can fluctuate, leading to a "spiky" or non-smooth time-averaged spectrum. Let's assume $A_k(t)$ is the amplitude (or related to amplitude) of mode k at frequency ω_k, and $\Phi_k(t)$ is its phase. The squaring suggests intensity, but the $\cos^2(\omega_k t + \cdots)$ part is problematic for a spectral density which should be a function of ω, not explicitly of t in this form.

Let's reinterpret: $I_{\mathrm{L}}(\omega, t)$ might be the *overall output intensity in time*, resulting from beating of modes. If it is the *spectral density*, then it should be

$$I_{\mathrm{L}}(\omega, t) = \sum_{\mathrm{k}} P_{\mathrm{k}}(t)\delta(\omega - \omega_{\mathrm{k}})$$

for ideal modes, where P_{k} is power of mode k.

Given the context of "peaks and holes" in spectral density, it likely means that the envelope of $P_{\mathrm{k}}(t)$ over k (i.e., over frequency) is not smooth.

"Time-averaged output:"

What we often measure with a spectrometer is the time-averaged spectral density.

Page 81:

The time-averaged spectral output $\langle I(\omega)\rangle$ is given by:

$$\langle I(\omega)\rangle = \frac{1}{T}\int_0^T \left[\sum_k A_{\mathrm{k}}(t)\cos^2\left(\omega_{\mathrm{k}} t + \Phi_{\mathrm{k}}(t)\right)\right] dt$$

"The time average of $\langle I(\omega)\rangle$ equals one over T, times the integral from zero to T, of the sum over k of $A_{\mathrm{k}}(t)$ times cosine squared of $\omega_{\mathrm{k}} t + \Phi_{\mathrm{k}}(t)$, dt."

If $A_{\mathrm{k}}(t)$ and $\Phi_{\mathrm{k}}(t)$ are slowly varying compared to $\frac{1}{\omega_{\mathrm{k}}}$ but perhaps fluctuate over the averaging time T, this average could still be non-smooth if the set of ω_{k} that are active changes, or their relative amplitudes A_{k} change.

If $A_{\mathrm{k}}(t)$ represents the slowly varying envelope of mode k at frequency ω_{k}, and $\cos^2$ represents its lineshape (unlikely), the issue is simpler: the A_{k} values themselves (the mode powers) may not form a smooth envelope.

* "For spectroscopic scanning, holes cause artificial features if the scanning time $< T$."

Imagine you are trying to measure an absorption spectrum by scanning the laser frequency across an absorption line. If the laser's output power spectrum $I_{\mathrm{L}}(\omega)$ has its own "holes" or "peaks" (i.e., it's not a flat baseline), and if these features in the laser spectrum fluctuate or are stable but sharp, they can be mistaken for absorption features of the sample, or they can distort the true absorption features.

This is particularly problematic if the time T_{scan} it takes to scan the laser over a spectral feature is shorter than the characteristic time T over which the laser's own spectral output averages to something smooth. Or, if the laser has fixed "holes" in its multimode output, these will appear as structure in any recorded spectrum.

* "Experimental remedy:"

How can we get a smoother effective laser spectrum for scanning?

* "Wobble the cavity length $d(t) = d_0 + \delta d \sin(2\pi f t)$ with $f > \frac{1}{T}$."

Here, d_0 is the mean cavity length, δd is a small modulation amplitude, and f is the modulation frequency. T would be the desired averaging time for the spectroscopic measurement. By rapidly modulating ("wobbling") the cavity length, all the axial mode frequencies $\nu_{\mathrm{q}} = \frac{qc}{2\, d(t)}$ are swept back and forth in frequency.

* "Rapidly modulates all $\omega_{\mathrm{k}} \Rightarrow$ smoother averaged spectrum."

If this wobbling is fast enough $\left(f > \frac{1}{T_{\text{measurement per point}}}\right)$, then each measurement point in the spectroscopic scan effectively averages over many different instantaneous sets of laser modes. This "smears out" the sharp peaks and holes in the laser's intrinsic multimode spectrum, leading to a more uniform, smoother effective output spectrum over the averaging time. This reduces baseline artifacts in the measured absorption spectrum. This technique is sometimes called "mode scrambling" or "dithering."

Page 82:

Slide 27:

Now we introduce a subtle but important effect: "Mode Pulling — Concept Introduction."

* "Even after mode selection, the exact frequency of a laser mode differs from the passive cavity resonance ν_r (nu sub arr) because of the frequency-dependent index $n(\nu)$ (en of nu)." Suppose we've managed to select a single axial and transverse mode. We might think its frequency is simply given by the passive cavity condition $\nu_\text{r} = \frac{qc}{2\,d}$.

However, this is not quite right. The active gain medium itself has a refractive index $n(\nu)$ that is frequency-dependent, especially near the gain line center due to anomalous dispersion (as we saw on pages 14-17).

The actual lasing frequency ν_a (nu sub ay, for active) will be determined by the round-trip phase condition including $n(\nu)$, i.e., $\nu_\text{a} = \frac{q c}{2 d^*(\nu_\text{a})}$, where $d^*(\nu_\text{a}) = d - L + n(\nu_\text{a})L$.

Because $n(\nu_\text{a})$ is generally not equal to 1 (the vacuum value), ν_a will be different from the passive cavity resonance ν_r (which assumes n=1 throughout).

* "This shift is called mode pulling." The presence of the dispersive gain medium "pulls" the lasing frequency away from where it would have been for an empty cavity.

* "Two intuitive viewpoints:"

1. "Cavity view: refractive index changes effective length d^* (dee star)." This is the view we've mostly used so far. The refractive index $n(\nu)$ of the gain medium changes the optical path length inside the cavity to $d^*(\nu)$. Since the resonance condition is $q\lambda = 2\,d^*$, a change in d^* (due

to $n(\nu)$ being different from 1 and varying with ν) will change the resonant wavelength and thus the frequency.

Page 83:

The second intuitive viewpoint for mode pulling:

2. "Gain medium view: phase of standing wave adjusted so that round-trip phase remains an integer multiple of π (pi) in presence of dispersion." (Should be integer multiple of 2π for phase, or for path length, integer multiple of λ).

The total round-trip phase shift $\phi(\nu)$ must equal $2\pi q$ for resonance.

$$\phi(\nu) = \frac{2\pi\nu}{c}\left[2\big((d-L) + n(\nu)L\big)\right]$$

(Using the one-way optical path $(d-L) + n(\nu)L$).

The term $n(\nu)L$ introduces a phase shift $\left(\frac{2\pi\nu}{c}\right)n(\nu)L$ within the gain medium (for one pass).

If $n(\nu)$ varies with frequency, the laser frequency ν must adjust itself slightly so that this total phase condition is met.

The frequency ν will shift from the passive cavity value ν_r to a new value ν_a such that the new phase $\phi(\nu_\text{a})$ (which includes $n(\nu_\text{a})$) is equal to $2\pi q$.

This adjustment of ν_a to maintain the phase condition in the presence of the dispersive $n(\nu)$ is mode pulling.

* "Quantitative derivation on next slide."

We will now derive a formula that quantifies this frequency shift.

Page 84:

Let's proceed with the "Mode Pulling — Mathematical Derivation Step-by-Step." (Slide 28)

* Passive cavity phase per round-trip:

Let ν_p (nu sub pee) be the resonant frequency of the passive cavity (what we called ν_r or $\nu_{q,vac}$ before). The round-trip phase Φ_p (Capital Phi sub pee) for this passive cavity is:

$$\Phi_\mathrm{p} = 2\pi\nu_\mathrm{p}\left(\frac{2\,d}{c}\right) = m\pi$$

Capital Phi sub pee equals two pi nu sub pee, times (two dee over cee), equals em pi.

Note: m here must be an even integer, $m = 2\,q$, for the phase to be $2\,q\pi$. If m is just an integer, it means $m\lambda/2 = d$ for path length. Let's assume $m\pi$ represents $2\,q\pi$, so m is effectively the axial mode number q if we use 2π for phase. Or, if m is q, then the phase shift is $q\pi$ for a single pass if node-node.

Let's stick to $\Phi = 2\pi q$ for round-trip phase for resonance.

So, $2\pi\nu_\mathrm{p}\left(\frac{2\,d}{c}\right) = 2\pi q$, which simplifies to $\nu_\mathrm{p}\left(\frac{2\,d}{c}\right) = q$, or $\nu_\mathrm{p} = \frac{qc}{2\,d}$. This is correct.

The slide uses $m\pi$. This would be correct if m is the number of full wavelengths in a round trip, then $m \cdot 2\pi$. If m is the number of half-wavelengths in a single trip d, then $m \cdot \pi$ for single trip phase. This can be confusing. Let's assume $m\pi$ here means "an integer multiple of π that satisfies resonance."

The standard condition is round trip phase = $2\pi q$. So, $2\pi\nu_\mathrm{p}\left(\frac{2\,d}{c}\right) = 2\pi q$.

* Active cavity phase:

Let ν_a (nu sub ay) be the actual lasing frequency in the active cavity.

The round-trip phase Φ_a (Capital Phi sub ay) includes the refractive index $n(\nu_a)$ of the gain medium (assuming it fills the cavity of length d for simplicity here, or d is d^*).

$$\Phi_a = 2\pi\nu_a\left(\frac{2\,d\,n(\nu_a)}{c}\right) = m\pi \quad \text{(again, } m\pi \text{ should be } 2\pi q\text{)}$$

Capital Phi sub ay equals two pi nu sub ay, times (two dee en of nu sub ay, over cee), equals em pi.

This formula implies the gain medium of index $n(\nu_a)$ fills the entire cavity length d. If it's of length L within d, then

$$\Phi_a = \frac{2\pi\nu_a}{c}\cdot 2[(d-L)+n(\nu_a)L].$$

The derivation will likely simplify this. For resonance, $\Phi_p = \Phi_a = 2\pi q$.

So, $\nu_p\,d = \nu_a[(d-L)+n(\nu_a)L]$ (after dividing by $2\pi(2/c)$ and q). This is one way to get ν_a.

Let's follow the slide's approach which seems to use a Taylor expansion, likely of $\Phi_a(\nu)$ around ν_p.

Page 85:

The slide proceeds by expanding $\Phi(\nu)$ (which is $\Phi_a(\nu)$) in a Taylor series around ν_p (the passive cavity resonance). This is a common technique when the shift $\nu_a - \nu_p$ is small. The phase for the active cavity is

$$\Phi_a(\nu) = \frac{4\pi\nu}{c}[(d-L)+n(\nu)L].$$

* "Expand $\Phi(\nu)$ in Taylor series around ν_p:" (The slide writes $\Phi(\nu)$ which refers to $\Phi_a(\nu)$) * $\Phi_a(\nu_a) \approx \Phi_a(\nu_p) + \left(\frac{d\Phi_a}{d\nu}\right)|_{\nu_p}(\nu_a - \nu_p)$ * The slide writes:

$$\Phi_a \approx \Phi_p + \left(\frac{\partial\Phi}{\partial\nu}\right)|_{\nu_p}(\nu_a - \nu_p) + [\Phi_a(\nu_a) - \Phi_p(\nu_a)] = m\pi$$

* This expansion is a bit unusual. Φ_p is $\Phi_\text{passive}(\nu_\text{p})$. The derivative term $\left(\frac{\partial \Phi}{\partial \nu}\right)|_{\nu_\text{p}}$ needs clarification on which Φ it is (active or passive).

* The term $[\Phi_\text{a}(\nu_\text{a}) - \Phi_\text{p}(\nu_\text{a})]$ looks like the difference in phase between active and passive cavities, evaluated at ν_a.

* This equation is effectively setting $\Phi_\text{a}(\nu_\text{a}) = m\pi$ (resonance for active cavity).

Let's assume a standard approximation for small pulling:

* The active cavity resonance ν_a must satisfy

$$\Phi_\text{active}(\nu_\text{a}) = 2\pi q.$$

* The passive cavity resonance ν_p satisfies

$$\Phi_\text{passive}(\nu_\text{p}) = 2\pi q.$$

* So $\Phi_\text{active}(\nu_\text{a}) = \Phi_\text{passive}(\nu_\text{p})$.

* $\Phi_\text{active}(\nu) = \frac{4\pi\nu}{c}[(d - L) + n(\nu)L]$.

* $\Phi_\text{passive}(\nu) = \frac{4\pi\nu}{c} d$.

* The slide's Φ_a seems to be the active phase, and Φ_p the passive phase at ν_p.

* "Setting $\Phi_\text{p} = m\pi$ and rearranging yields:" If

$$\Phi_\text{p} = \Phi_\text{passive}(\nu_\text{p}) = m\pi \quad \text{(i.e. } 2\pi q\text{)},$$

and we want $\Phi_\text{active}(\nu_\text{a}) = m\pi$, then the equation implies:

$$\left(\frac{\partial \Phi}{\partial \nu}\right)|_{\nu_\text{p}}(\nu_\text{a} - \nu_\text{p}) + [\Phi_\text{a}(\nu_\text{a}) - \Phi_\text{p}(\nu_\text{a})] = 0.$$

* Partial Φ partial ν evaluated at ν_p, times $\left(\nu_\mathrm{a} - \nu_\mathrm{p}\right)$, plus open bracket $\Phi_\mathrm{a}(\nu_\mathrm{a}) - \Phi_\mathrm{p}(\nu_\mathrm{a})$ close bracket, equals zero.

* The slide states this is "identical to Eq. (5.72)" from the textbook. This equation relates the frequency shift $\nu_\mathrm{a} - \nu_\mathrm{p}$ to the difference in phase response of the active and passive cavities.

* The term $\Phi_\mathrm{a}(\nu_\mathrm{a}) - \Phi_\mathrm{p}(\nu_\mathrm{a})$ can be written as

$$\frac{4\pi\nu_\mathrm{a}}{c}\, L\,(n(\nu_\mathrm{a}) - 1).$$

This is the extra phase due to the medium of length L.

* The derivative $\left(\frac{\partial\Phi}{\partial\nu}\right)|_{\nu_\mathrm{p}}$ is related to the group delay or cavity storage time. For a passive cavity, $\Phi_\mathrm{p}(\nu) = \frac{4\pi\nu d}{c}$, so

$$\frac{\partial\Phi_\mathrm{p}}{\partial\nu} = \frac{4\pi d}{c} = 2\pi\left(\frac{2\,d}{c}\right) = 2\pi\,\tau_\mathrm{rt},$$

where τ_rt is the round trip time.

This formulation is a bit dense without the context of Eq. (5.72). However, the goal is to find $\nu_\mathrm{a} - \nu_\mathrm{p}$.

* "Need dispersion relation linking $n(\nu)$ to gain coefficient $\alpha(\nu)$:" To solve this, we need an expression for $n(\nu)$, specifically how it varies around the gain line because of the gain $\alpha(\nu)$ (or $g(\nu) = -\alpha(\nu)$). This is given by the Kramers-Kronig relations, often approximated for a Lorentzian gain line.

Page 86:

The dispersion relation linking $n(\nu)$ to the absorption coefficient $\alpha(\nu)$ (where gain corresponds to negative $\alpha(\nu)$) near a resonance ν_0 is given (often for a Lorentzian line) as:

$$n(\nu) = 1 + \left(\frac{\nu_0 - \nu}{\Delta\nu_\mathrm{m}}\right)\cdot\left(\frac{c}{2\pi\nu}\right)\alpha(\nu)$$

Let's define the terms:

* ν_0 (nu naught): "centre frequency of gain line."

$\Delta\nu_\mathrm{m}$ (Delta nu sub em): This is related to the width of the gain line. The slide defines it as $\Delta\nu_\mathrm{m} = \frac{\gamma_\mathrm{m}}{2\pi}$: "homogeneous half-width." Here γ_m would be the HWHM in angular frequency units if $\Delta\nu_\mathrm{m}$ is HWHM in Hz. Or, if γ_m is the full homogeneous linewidth (angular), then $\Delta\nu_\mathrm{m}$ is the FWHM (Hz) if the relation is $\Delta\nu_\mathrm{m} = \frac{\gamma_\mathrm{m}}{2\pi}$. Let's assume $\Delta\nu_\mathrm{m}$ is the HWHM in Hz of the gain line profile*. (Often $\Delta\nu_\mathrm{L}$ or $\Delta\nu_\mathrm{h}$ used for this).

* c is speed of light. $\alpha(\nu)$ is the absorption coefficient (negative for gain). The $2\pi\nu$ in the denominator is sometimes approximated as $2\pi\nu_0$ if the frequency range is narrow. This formula gives the change in refractive index due to the resonant gain/absorption feature $\alpha(\nu)$. The 1 is the background refractive index (vacuum).

* "After algebra (see next slide) obtain the mode-pulled frequency." By substituting this expression for $n(\nu)$ into the phase condition equation (implicitly, through Eq. 5.72 from the previous slide), and performing some algebraic manipulations (which can be quite involved), one can derive an expression for the actual lasing frequency ν_a.

Page 87:

This slide presents the "Final Formula and Limiting Cases" for Mode Pulling. (Slide 29)

* "Derived expression:"

The actual lasing frequency ν_a is given by:

$$\nu_\mathrm{a} = \frac{\nu_\mathrm{r}\Delta\nu_\mathrm{m} + \nu_0\Delta\nu_\mathrm{r}}{\Delta\nu_\mathrm{m} + \Delta\nu_\mathrm{r}}$$

Let's define the terms in this very important formula: "where" (Definitions continue on the next page)

ν_{a}: is the actual frequency of the laser mode (the mode-pulled frequency).

ν_{r}: is the passive cavity resonance frequency (what we called ν_{p} or $\nu_{q,vac}$ earlier). This is the frequency the mode *would* have if there were no gain medium dispersion.

ν_0: is the center frequency of the gain line profile.

$\Delta\nu_{\mathrm{m}}$: is the HWHM (Half Width at Half Maximum) of the *gain line* (the "m" might stand for medium or molecular transition). This was on the previous slide.

$\Delta\nu_{\mathrm{r}}$: is the HWHM of the *passive cavity resonance* (the "r" for resonator).

This formula is beautifully intuitive: the actual lasing frequency ν_{a} is a weighted average of the passive cavity frequency ν_{r} and the gain line center frequency ν_0. The weighting factors are the half-widths of the *other* feature: ν_{r} is weighted by $\Delta\nu_{\mathrm{m}}$, and ν_0 is weighted by $\Delta\nu_{\mathrm{r}}$.

Page 88:

Let's continue with the definitions and interpretation for the mode pulling formula:

$$\nu_{\mathrm{a}} = \frac{\nu_{\mathrm{r}}\,\Delta\nu_{\mathrm{m}} + \nu_0\,\Delta\nu_{\mathrm{r}}}{\Delta\nu_{\mathrm{m}} + \Delta\nu_{\mathrm{r}}}$$

* $\Delta\nu_{\mathrm{r}} = \frac{c\,\gamma_{\mathrm{loss}}}{4\pi d}$ (Corrected formula likely from context, slide has $\frac{c\gamma}{4\pi d}$ without specifying γ). The slide says $\Delta\nu_{\mathrm{r}} = \frac{c\gamma}{4\pi d'}$. This γ seems to be related to cavity losses. More commonly, $\Delta\nu_{\mathrm{r}}$ (the HWHM of the passive cavity resonance) is related to the FSR and finesse F_{p}:

$$\Delta\nu_{\mathrm{r}} = \frac{\mathrm{FWHM}_{\mathrm{cavity}}}{2} = \frac{\delta\nu/F_{\mathrm{p}}}{2} = \frac{c/\left(2\,dF_{\mathrm{p}}\right)}{2} = \frac{c}{4\,dF_{\mathrm{p}}}$$

If the γ on the slide here refers to the round-trip intensity loss $(1-R_{\text{eff}})$, and $F_{\text{p}} \approx \pi/(1-R_{\text{eff}})$, then $\Delta\nu_{\text{r}} \approx \frac{c(1-R_{\text{eff}})}{4\pi d}$. So if $\gamma = (1-R_{\text{eff}})$, this is consistent. The definition given is: "passive cavity resonance half-width (determined by total losses γ)." This means $\Delta\nu_{\text{r}}$ is the HWHM of the cavity resonance. Its narrowness depends on cavity losses: low loss means high finesse and small $\Delta\nu_{\text{r}}$.

* "Interpretation:" * "Weighted average of cavity resonance ν_{r} and gain centre ν_0." This is clear from the form of the equation. * "Weighting factors are respective half-widths." This is the key insight. Whichever feature (cavity resonance or gain line) is "sharper" (has a smaller half-width) has more influence on the other. No, that's not right. The weighting of ν_{r} is $\frac{\Delta\nu_{\text{m}}}{\Delta\nu_{\text{m}}+\Delta\nu_{\text{r}}}$ and the weighting of ν_0 is $\frac{\Delta\nu_{\text{r}}}{\Delta\nu_{\text{m}}+\Delta\nu_{\text{r}}}$. So, if $\Delta\nu_{\text{m}}$ (gain linewidth) is much larger than $\Delta\nu_{\text{r}}$ (cavity linewidth), i.e., $\Delta\nu_{\text{m}} \gg \Delta\nu_{\text{r}}$, then $\nu_{\text{a}} \approx \frac{\nu_{\text{r}}\,\Delta\nu_{\text{m}}}{\Delta\nu_{\text{m}}} = \nu_{\text{r}}$. This means if the gain line is very broad and the cavity resonance is very sharp, the lasing frequency ν_{a} is pulled very little and stays close to ν_{r}. The sharp cavity dominates. Conversely, if $\Delta\nu_{\text{r}} \gg \Delta\nu_{\text{m}}$ (cavity resonance is very broad, gain line is very sharp), then $\nu_{\text{a}} \approx \frac{\nu_0\,\Delta\nu_{\text{r}}}{\Delta\nu_{\text{r}}} = \nu_0$. The lasing frequency is pulled strongly towards the gain line center ν_0. The sharp gain line dominates. This makes sense: the resulting frequency is pulled towards the "sharper" of the two original frequencies (passive cavity or gain line center). (Wait, the weighting is by the *other's* width. If $\Delta\nu_{\text{m}}$ is large, it means ν_{r} has a large weight. So, if gain line is broad, frequency is pulled *towards* ν_{r}. If cavity line is broad (small finesse), frequency is pulled *towards* ν_0. This is it!) So, the laser frequency ν_{a} will be closer to ν_{r} if $\Delta\nu_{\text{m}}$ is large (broad gain line). It will be closer to ν_0 if $\Delta\nu_{\text{r}}$ is large (broad cavity resonance). This means ν_{a} is pulled towards the component (ν_{r} or ν_0) whose associated feature ($\Delta\nu_{\text{r}}$ or $\Delta\nu_{\text{m}}$ respectively) is NARROWER. Yes, this is correct. The formula effectively says the output frequency ν_{a} is "pulled" from ν_{r} towards ν_0 by an

amount proportional to $\Delta\nu_{\mathrm{r}}$ and inversely proportional to $\Delta\nu_{\mathrm{m}}$ (for small pulls).

* "Gas lasers typical numbers:" * $\Delta\nu_{\mathrm{r}} \sim 1\,$MHz. (Delta nu sub arr is around one megahertz). This is the HWHM of a typical good quality passive cavity resonance. FWHM would be $\sim 2\,$MHz. (This implies a finesse. If FSR is $150\,$MHz, FWHM 2MHz means Finesse = $150/2 = 75$. This is reasonable.)

Page 89:

Continuing with typical numbers for gas lasers:

- $\Delta\nu_{\mathrm{m}} \sim 100\,$MHz. (Delta nu sub em is around one hundred megahertz). This is the HWHM of the homogeneous gain line profile. (This could be, for instance, the pressure-broadened linewidth in a He-Ne laser, not the much larger Doppler width). The mode pulling phenomenon is driven by the dispersion associated with the part of the gain that is actually being "used" by the mode, which can be related to the homogeneous packet.

- "Simplified form for $\Delta\nu_{\mathrm{r}} \ll \Delta\nu_{\mathrm{m}}$:" (Delta nu sub arr is much less than Delta nu sub em). This is the common case for gas lasers: sharp cavity resonance ($\Delta\nu_{\mathrm{r}} \sim 1\,$MHz), broader homogeneous gain line ($\Delta\nu_{\mathrm{m}} \sim 100\,$MHz). In this limit, the formula

$$\nu_{\mathrm{a}} = \frac{\nu_{\mathrm{r}}\Delta\nu_{\mathrm{m}} + \nu_0\Delta\nu_{\mathrm{r}}}{\Delta\nu_{\mathrm{m}} + \Delta\nu_{\mathrm{r}}}$$

can be approximated.

Divide numerator and denominator by $\Delta\nu_{\mathrm{m}}$:

$$\nu_{\mathrm{a}} = \frac{\nu_{\mathrm{r}} + \nu_0\left(\frac{\Delta\nu_{\mathrm{r}}}{\Delta\nu_{\mathrm{m}}}\right)}{1 + \frac{\Delta\nu_{\mathrm{r}}}{\Delta\nu_{\mathrm{m}}}}$$

Using binomial expansion

$$(1 + x)^{-1} \approx 1 - x$$

for small $x = \frac{\Delta\nu_{\mathrm{r}}}{\Delta\nu_{\mathrm{m}}}$:

$$\nu_{\mathrm{a}} \approx \left(\nu_{\mathrm{r}} + \nu_0\left(\frac{\Delta\nu_{\mathrm{r}}}{\Delta\nu_{\mathrm{m}}}\right)\right)\left(1 - \frac{\Delta\nu_{\mathrm{r}}}{\Delta\nu_{\mathrm{m}}}\right)$$

$$\nu_{\mathrm{a}} \approx \nu_{\mathrm{r}} - \nu_{\mathrm{r}}\left(\frac{\Delta\nu_{\mathrm{r}}}{\Delta\nu_{\mathrm{m}}}\right) + \nu_0\left(\frac{\Delta\nu_{\mathrm{r}}}{\Delta\nu_{\mathrm{m}}}\right) - \nu_0\left(\frac{\Delta\nu_{\mathrm{r}}}{\Delta\nu_{\mathrm{m}}}\right)^2$$

Ignoring the squared term:

$$\nu_{\mathrm{a}} \approx \nu_{\mathrm{r}} + (\nu_0 - \nu_{\mathrm{r}})\left(\frac{\Delta\nu_{\mathrm{r}}}{\Delta\nu_{\mathrm{m}}}\right)$$

This is the formula given on the slide:

$$\nu_{\mathrm{a}} \approx \nu_{\mathrm{r}} + \left(\frac{\Delta\nu_{\mathrm{r}}}{\Delta\nu_{\mathrm{m}}}\right)(\nu_0 - \nu_{\mathrm{r}})$$

"nu sub ay is approximately nu sub arr, plus (Delta nu sub arr over Delta nu sub em) times (nu naught minus nu sub arr)."

- "Mode pulled only slightly towards ν_0." Since $\Delta\nu_{\mathrm{r}} \ll \Delta\nu_{\mathrm{m}}$, the ratio $\frac{\Delta\nu_{\mathrm{r}}}{\Delta\nu_{\mathrm{m}}}$ is small (e.g., $\frac{1\ \mathrm{MHz}}{100\ \mathrm{MHz}} = 0.01$). So the shift $\nu_{\mathrm{a}} - \nu_{\mathrm{r}}$ is only a small fraction of the detuning $\nu_0 - \nu_{\mathrm{r}}$. This means ν_{a} stays very close to ν_{r}. The pulling is weak, which is expected if the cavity is much sharper than the gain line feature causing the dispersion.
- "Pulling magnitude proportional to detuning ($\nu_0 - \nu_{\mathrm{r}}$)." The further the passive cavity mode ν_{r} is from the gain line center ν_0, the larger the pull $\nu_{\mathrm{a}} - \nu_{\mathrm{r}}$ will be (for a fixed ratio of linewidths). The pulling is zero if $\nu_{\mathrm{r}} = \nu_0$ (cavity tuned to line center).

Page 90:

Let's consider the "Practical Consequences of Mode Pulling." (Slide 30)

* "In line-centre operation ($\nu_\mathrm{r} \approx \nu_0$) pulling is negligible." This follows directly from the formula

$$\nu_\mathrm{a} - \nu_\mathrm{r} \approx \left(\frac{\Delta\nu_\mathrm{r}}{\Delta\nu_\mathrm{m}}\right)(\nu_0 - \nu_\mathrm{r})$$

If $\nu_\mathrm{r} = \nu_0$, then $(\nu_0 - \nu_\mathrm{r}) = 0$, so $\nu_\mathrm{a} = \nu_\mathrm{r}$. There is no pulling. This is an important operating point for minimizing pulling effects.

* "Near slope of gain curve:" (This refers to when ν_r is on the slope of the gain curve, i.e., detuned from ν_0). * "Small cavity length drifts ⇒ large frequency drifts partially cancelled by pulling." This is a subtle but beneficial effect. Suppose the cavity length d changes by a small amount Δd (e.g., due to thermal expansion or vibration). This would cause the passive cavity frequency

$$\nu_\mathrm{r} = \frac{qc}{2d}$$

to shift by

$$\Delta\nu_{r,\mathrm{drift}} = -\left(\frac{qc}{2d^2}\right)\Delta d = -\frac{\nu_\mathrm{r}}{d}\Delta d$$

This can be a large frequency drift. However, as ν_r drifts, the detuning $(\nu_0 - \nu_\mathrm{r})$ changes. This change in detuning causes the mode pulling $(\nu_\mathrm{a} - \nu_\mathrm{r})$ to change in a way that partially counteracts the initial drift of ν_r. The pulling term shifts ν_a *towards* ν_0. So, if ν_r drifts away from ν_0, the pulling effect increases, pulling ν_a back towards ν_0 more strongly. If ν_r drifts towards ν_0, the pulling effect weakens. This results in the actual lasing frequency ν_a being somewhat more stable against cavity length drifts than ν_r would be on its own. The gain medium effectively provides a kind of "restoring force" that pulls ν_a back towards ν_0. This is sometimes called "frequency pulling compensation" or "self-correction." The degree of cancellation depends on the pulling factor

$$\frac{\Delta\nu_\mathrm{r}}{\Delta\nu_\mathrm{m}}.$$

* "Stabilization schemes exploit or compensate for this." Understanding mode pulling is crucial for designing laser frequency stabilization systems. If you are stabilizing the laser by locking ν_r to an external reference (e.g., another stable cavity), you need to be aware that ν_a will still be affected by pulling if ν_0 drifts or if the pulling factor changes. Some schemes might try to lock ν_a directly to an atomic reference, which bypasses issues with ν_r drift and pulling relative to ν_r. Other schemes might actively control ν_r to keep ν_a stable, using knowledge of the pulling effect. For example, if ν_0 is known and stable, one can lock ν_r to a specific offset from ν_0 to achieve a desired ν_a.

* "For precision spectroscopy:" This is where mode pulling becomes extremely important.

Page 91:

Continuing with practical consequences for precision spectroscopy:

* "Mode pulling introduces systematic error if laser is locked to cavity rather than atomic reference." This is a major concern. Many laser stabilization schemes involve locking the laser's cavity length d (and thus ν_r) to a high-finesse, stable but passive reference cavity (e.g., made of ultra-low expansion glass, kept in vacuum and temperature controlled). This can make ν_r very stable. However, the actual lasing frequency ν_a is what interacts with the atoms in a spectroscopy experiment. If $\nu_\mathrm{a} = \nu_\mathrm{r} +$ PullingTerm, and if the PullingTerm changes (e.g., due to drifts in the gain line center ν_0, or changes in laser power affecting $\Delta\nu_\mathrm{m}$ or $\Delta\nu_\mathrm{r}$ through saturation effects), then ν_a will drift even if ν_r is perfectly stable. This drift in ν_a would appear as a systematic error in frequency measurements if one assumes $\nu_\mathrm{a} = \nu_\mathrm{r}$. Locking directly to an "atomic reference" (e.g., using saturated absorption spectroscopy to lock ν_a to the center of an

unperturbed atomic transition) avoids this particular systematic error, because you are then directly controlling ν_a.

* "Correction requires knowledge of $\Delta\nu_m$, $\Delta\nu_r$, and detuning." If one must lock to a passive cavity and needs to know ν_a very accurately, then one has to measure or model the mode pulling effect. This requires knowing all the parameters in the mode pulling formula: ν_r (from the lock to the reference cavity), ν_0 (the gain line center, which might drift), $\Delta\nu_m$ (gain line HWHM, might depend on operating conditions), and $\Delta\nu_r$ (cavity HWHM, might also vary if intracavity losses change). Accurately determining all these to correct for mode pulling can be very challenging.

* "[IMAGE REQUIRED: Plot showing cavity resonance (vertical lines) and gain curve; arrows indicating mode shift towards centre as gain slope steepens.]" This describes a conceptual diagram: Imagine a broad gain curve (like a Gaussian). Underneath it, show several sharp vertical lines representing the passive cavity resonant frequencies ν_r. Then, for each ν_r, show an arrow pointing from ν_r towards the center of the gain curve (ν_0). The length of this arrow would represent the magnitude of the pull ($\nu_a - \nu_r$). The arrows should be longer for ν_r values further from ν_0 (on the slopes of the gain curve) because the detuning ($\nu_0 - \nu_r$) is larger there. The phrase "as gain slope steepens" is a bit ambiguous. The pulling is largest where the detuning ($\nu_0 - \nu_r$) is large, i.e., on the flanks of the gain curve. The slope of the *refractive index curve* $n(\nu)$ is what directly drives the pulling, and this slope $\frac{dn}{d\nu}$ is related to the gain curve $\alpha(\nu)$ via Kramers-Kronig. Typically, $\frac{dn}{d\nu}$ is largest near the center of a resonance for a gain line. The pulling magnitude $\nu_a - \nu_r$ is proportional to $(\nu_0 - \nu_r) \cdot \frac{\Delta\nu_r}{\Delta\nu_m}$. It is this detuning factor that is key.

Page 92:

We now reach a "Summary of Key Quantitative Relationships." This is Slide 31. This slide recaps some of the most important formulas we've encountered for describing laser spectra.

* "Passive cavity FSR (Free Spectral Range):"

$$\delta\nu = \frac{c}{2\,d}$$

"delta nu equals cee divided by two dee." This gives the constant spacing between adjacent axial modes of an empty (passive) cavity of length d.

"Round-trip gain factor:" (This refers to the round-trip amplitude* gain factor G)

$$G(\nu) = \exp[-2\alpha(\nu)L - \gamma(\nu)]$$

"Capital Gee of nu equals exponential of, open square bracket, minus two alpha of nu times capital Ell, minus gamma of nu, close square bracket." Here, $\alpha(\nu)$ is the absorption coefficient of the gain medium (negative for gain) of length L. $\gamma(\nu)$ is the total passive logarithmic amplitude loss per round-trip. For lasing, $G(\nu)$ must be ≥ 1.

* "Active cavity linewidth:" (This is the FWHM of a lasing mode, $\Delta\nu_{\mathrm{a}}$, approaching threshold) The formula given here is slightly different from page 42, missing a 2π factor, or $\delta\nu$ is FSR. The formula on page 42 was

$$\Delta\nu_{\mathrm{a}} = \delta\nu \cdot \left(\frac{1 - G(\nu)}{2\pi\sqrt{G(\nu)}}\right)$$

where $\delta\nu$ was FSR. The equation is not fully written on this summary slide, but it's implied by the next slide. Let's look at page 93 for the full formula.

Page 93:

Continuing the summary of key relationships, the active cavity linewidth formula is given at the top of this page:

$$\Delta\nu_{\mathrm{a}} = \delta\nu \cdot \frac{1 - G}{2\pi\sqrt{G}}$$

"Delta nu sub ay equals delta nu times, fraction: (one minus capital Gee) divided by (two pi times square root of capital Gee)."

Here, $\delta\nu$ is the FSR $\left(\frac{c}{2\,d}\right)$. G is the round-trip amplitude gain $G(\nu)$. As $G \to 1^{-}$, $\Delta\nu_{\mathrm{a}} \to 0$. This is the Schawlow-Townes narrowing before considering quantum noise.

* "Saturated gain coefficients:" How the gain (or absorption coefficient α) is reduced by the presence of laser intensity.

* "Homogeneous:" For a homogeneously broadened medium.

$$\alpha_{\mathrm{s}} = \frac{\alpha_0}{1 + I/I_{\mathrm{s}}}$$

"alpha sub ess equals alpha sub zero, divided by (one plus capital Eye over capital Eye sub ess)." Here, α_0 is the small-signal absorption coefficient, α_{s} is the saturated absorption coefficient, I is the intensity, and I_{s} is the saturation intensity. If α_0 is for gain (negative), then α_{s} becomes less negative.

* "Inhomogeneous:" For an inhomogeneously broadened medium.

Page 94:

The saturated gain coefficient for an inhomogeneously broadened medium:

$$\alpha_{\mathrm{s}} = \frac{\alpha_0}{\sqrt{1 + \frac{I}{I_{\mathrm{s}}}}}$$

Again, α_{s} saturates more slowly (denominator is larger, so $|\alpha_{\mathrm{s}}|$ is larger, closer to $|\alpha_0|$) compared to the homogeneous case for the same $\frac{I}{I_{\mathrm{s}}}$.

* Mode pulling: The formula for the actual lasing frequency ν_{a}:

$$\nu_a = \frac{\nu_r \Delta\nu_m + \nu_0 \Delta\nu_r}{\Delta\nu_m + \Delta\nu_r}$$

Where ν_r is passive cavity resonance, ν_0 is gain line center, $\Delta\nu_m$ is HWHM of gain line, and $\Delta\nu_r$ is HWHM of passive cavity resonance.

* A final crucial comment: "Each equation requires understanding of the definitions and units provided earlier; together they form the quantitative backbone for predicting and controlling laser spectra." This is absolutely true. These equations are not just abstract symbols; each term has a physical meaning, units, and depends on specific characteristics of the laser system. Mastering these relationships is essential for anyone who wants to design lasers, understand their output, diagnose problems, or use them for high-precision applications like laser spectroscopy. You need to know what each symbol represents, how it's typically measured or calculated, and how it influences the others. This set of formulas, built up step-by-step, allows us to move from basic principles to a fairly sophisticated understanding of why a laser emits the specific frequencies it does, how narrow those frequencies can be, and how they might shift or change.

Page 95

Slide 32

Alright, let's move to our "Concluding Remarks and Next Steps" for this chapter on the spectral characteristics of laser emission. This is Slide 32.

* The first key takeaway is a crucial one: "The frequency composition of laser emission is not a simple direct mirror of the gain profile; it emerges from a rich interplay of:" This is perhaps the most important conceptual summary. One might naively think that if you have a gain medium with a certain spectral gain curve, the laser will just emit light across that entire curve. But that's not the case. The actual spectrum is far more structured

and selective, resulting from a complex interplay of several factors that we've discussed. These factors include:

* "Cavity resonance conditions." The optical resonator only supports discrete frequencies (axial and transverse modes) where standing waves can form. This imposes the first level of selection. The laser can only operate at or very near these allowed cavity frequencies. * "Gain saturation (homogeneous vs inhomogeneous)." Once lasing begins, the intensity builds up and saturates the gain. How this saturation occurs – whether uniformly across the gain profile (homogeneous) or by burning selective spectral holes (inhomogeneous) – dramatically affects which modes can lase and how they compete with each other. * "Spatial hole burning and transverse-mode structure." In linear cavities, the standing wave pattern leads to spatial variations in saturation, which can allow multiple modes (even in a homogeneous medium) to coexist by drawing gain from different regions. The presence of higher-order transverse modes (TEM_{mn}) also adds complexity to the spectrum, as these modes have their own distinct frequencies. * "Dispersive frequency shifts (mode pulling)." The refractive index of the gain medium itself is frequency-dependent, especially near the gain line. This causes the actual lasing frequencies to be "pulled" away from the passive cavity resonances, towards the center of the gain line. The magnitude of this pulling depends on the relative widths of the cavity resonance and the gain line.

It's the combination of all these effects – the discrete cavity modes, the shape of the gain, how the gain saturates, where it saturates spatially, and how the medium's dispersion shifts frequencies – that ultimately determines the detailed spectral output of any given laser.

Page 96:

Continuing with our concluding remarks:

"Experimental control levers:"

Given this complex interplay, what tools do we have as experimentalists to influence and control the laser's spectral output? We have several "levers":

1. "Adjust cavity length d and geometry." Changing the cavity length d directly changes the FSR ($\frac{c}{2d}$) and shifts the entire comb of axial mode frequencies. Fine-tuning d is essential for selecting a specific mode or tuning the laser. Cavity geometry (mirror curvatures, stability parameters) affects the transverse mode spacing and stability, and can be used to favor or suppress certain modes (like TEM_{00}).

2. "Shape gain profile via temperature, pressure, and pumping." The gain profile $\alpha(\nu)$ (or $g(\nu)$) is not always fixed. For some gain media, temperature can affect the center frequency and width. For gas lasers, pressure influences collisional broadening and thus $\Delta\nu_{\text{hom}}$. Pumping level directly controls the magnitude of the small-signal gain, which determines how far above threshold the laser operates and thus the width of the lasing band $\Delta\nu_{\text{lase}}$.

3. "Insert selective elements (etalons, gratings, birefringent filters)." As we've discussed, particularly for broad-gain media or when single-frequency operation is critical, intracavity elements are used to introduce frequency-dependent losses. These "tune" the net gain curve, allowing us to pick out a single axial and transverse mode and often to tune its frequency over a portion of the gain bandwidth.

"Subsequent chapters/slides will extend these concepts to:"

This section has laid the groundwork. Where do we go from here? We will build upon these ideas to explore even more advanced topics in laser spectroscopy:

"Linewidth narrowing mechanisms (Schawlow-Townes limit)." We briefly mentioned that there's a fundamental quantum limit to laser linewidth. We will delve into the physics of this limit and explore techniques and cavity designs that aim to approach it, leading to ultra-narrow linewidth lasers.

"Active frequency stabilization techniques." Beyond passive stability and basic mode selection, achieving very high frequency stability (e.g., parts in 10^{15} or better) requires active feedback systems. These involve locking the laser frequency to an external stable reference, such as an atomic transition or a high-finesse optical cavity, using techniques like the Pound-Drever-Hall method.

Page 97:

And further extensions of these concepts will include:

"Tunable and mode-locked laser systems."

We've touched on tunability achieved by intracavity elements. We'll explore various designs for broadly tunable lasers, which are workhorses for spectroscopy.

"Mode-locked laser systems" are a different but related topic. By forcing many axial modes of a laser to oscillate with a fixed phase relationship, one can generate trains of ultra-short (femtosecond or picosecond) pulses of light. Understanding the axial mode structure $\delta\nu = \frac{c}{2d}$ is fundamental to mode-locking, as the pulse repetition rate is typically equal to the FSR or a sub-harmonic. Mode-locked lasers have their own unique spectral characteristics (a broad comb of phase-locked frequencies).

Finally, "Mastery of today's material equips you to design, diagnose, and optimize laser sources for high-resolution spectroscopy and other precision applications."

This is the ultimate goal. The detailed understanding of how active media and optical resonators interact to define the laser's spectral output is not just academic. It is intensely practical.

If you are performing high-resolution spectroscopy, you need to know if your laser is single-mode, what its linewidth is, how stable its frequency is, and what might be causing unwanted spectral features or drifts.

If you are designing a laser system for a specific application, you need to choose the right gain medium, design an appropriate resonator, and incorporate the necessary control elements to achieve the desired spectral performance (e.g., specific wavelength, narrow linewidth, tunability, power level).

If your laser isn't behaving as expected (e.g., it's multimode when you want single-mode, or its frequency is unstable), the concepts covered here – gain saturation, spatial and spectral hole burning, mode pulling, competition, cavity resonances – provide the framework for diagnosing the problem and finding a solution.

Many "other precision applications," from optical clocks and metrology to coherent communication and quantum information processing, rely heavily on lasers with exquisitely controlled spectral properties.

So, the material we've covered in this chapter, while detailed, forms an essential toolkit for any physicist or engineer working seriously with lasers. It's the foundation upon which many advanced laser techniques and applications are built.

This concludes our discussion for Chapter 5.3. Thank you.