

Passive mode locking of thin-disk lasers: effects of spatial hole burning

R. Paschotta¹, J. Aus der Au¹, G.J. Spühler¹, S. Erhard², A. Giesen², U. Keller¹

¹Ultrafast Laser Physics, Institute of Quantum Electronics, Swiss Federal Institute of Technology, ETH Hönggerberg HPT, 8093 Zürich, Switzerland

²Institut für Strahlwerkzeuge, Universität Stuttgart, Pfaffenwaldring 43, 70569 Stuttgart, Germany

Received: 4 September 2000/Published online: 10 January 2001 – © Springer-Verlag 2001

Abstract. We recently demonstrated that passive mode locking of a thin-disk Yb:YAG laser is possible and that this concept leads to sources of femtosecond pulses with very high average power. Here we discuss in detail the effect of spatial hole burning on the mode-locking behavior of such lasers. We have developed an efficient numerical model and arrive at quantitative stability criteria which agree well with experimental data. The main result is that stable soliton mode locking can in general be obtained only in a certain range of pulse durations. We use our model to investigate the influence of various cavity parameters and the situation for different gain media. We also consider several methods to reduce the effect of spatial hole burning in order to expand the range of possible pulse durations.

PACS: 42.60.Fc; 42.55.Rz; 42.55.Xi

Until recently, the average output power of passively mode-locked all-solid-state lasers was limited to below 1 W. Particularly in the sub-picosecond domain, the usually quite poor thermal properties of the available broadband gain media precluded the generation of higher powers. Only a short while ago we have shown [1] that thin-disk Yb:YAG lasers [2], which had been demonstrated to generate up to ≈ 100 W of nearly diffraction-limited continuous-wave power [3], can be passively mode-locked using a semiconductor saturable absorber mirror (SESAM) [4–6]. In this first demonstration, we obtained 16 W of average power in 0.7-ps pulses. Furthermore, we have shown that the concept is power-scalable so that substantially higher output powers should be possible in the near future. As mentioned already in [1], spatial hole burning (SHB) has a strong effect on the mode-locking behavior because it leads to inhomogeneous gain saturation. In this paper, we discuss this effect in detail.

The effect of SHB on the mode-locking behavior of lasers (although not specifically of thin-disk lasers) has already been discussed in other papers, in particular in [7, 8]. It was found that SHB can facilitate the generation of shorter pulses, al-

though the time–bandwidth product may be deteriorated. It was not anticipated that stable soliton mode locking [9] of a laser with SHB may be possible only in a narrow range of pulse durations, as we observed in our experiments with a thin-disk laser [1]. We can quantitatively explain this observation with a new model, which is at the same time simpler and numerically more efficient than the model used in [8], because it makes use of the fact that the pulse spectrum is more or less fixed in a soliton mode-locked laser where the pulse is shaped mainly by group-delay dispersion and self-phase modulation. We derive quantitative stability criteria, based on the calculation of the effective gain of a continuous-wave (cw) background (often called the ‘continuum’ in soliton theory) and of double pulses, which can compete with the (desired) single-soliton pulse in the laser cavity.

The paper is organized as follows. First, in Sect. 1 we briefly describe the setup of a thin-disk laser and present an efficient numerical method to calculate the spectral shape of the saturated gain. In Sect. 2 we discuss how this spectral shape can favor the formation of double pulses and thus destabilize the mode-locking process. We will show that stable mode locking is possible only if the parameters of the saturable absorber are properly chosen. In particular, for a given degree of absorber saturation there is a lower limit for the modulation depth (maximum reflectivity change), set by the tendency for unstable double pulsing, as well as a higher limit set by Q-switching instabilities [10, 11]. As both the lower and the upper limits depend on the soliton-pulse duration, the cavity dispersion must be appropriately chosen. We will mainly concentrate on a slightly modified version of the laser described in [1], but also discuss possible changes of the design and their effect on the mode-locking behavior. In Sect. 3 we discuss the effect of SHB in thin-disk lasers with other gain media, both those with a smaller and with a larger amplification bandwidth, and show that this leads into regimes with qualitatively different behavior. Finally, in Sect. 4 we consider various methods to reduce the effect of SHB in mode-locked thin-disk lasers, which can be desirable to obtain mode locking in a larger range of pulse durations.

1 Gain in a thin-disk laser

First we briefly discuss the geometry of a thin-disk laser head (Fig. 1), which has been described in more detail in [2]. The Yb:YAG disk with a thickness in the order of 0.2 mm has a highly reflecting coating on the side which is attached to the cooling finger. Thus the laser beam always generates a standing-wave pattern in the disk, regardless of the type of laser cavity (ring or standing-wave), and SHB unavoidably occurs. The outer disk surface has an anti-reflection (AR) coating. For mode-locked operation, the disk surfaces are wedged by $\approx 0.1^\circ$ in order to definitely eliminate the disturbing effect of residual reflections. An optical arrangement [12] consisting of a parabolic mirror and three roof prisms is used to obtain multiple passes of the pump radiation through the thin disk, in order to obtain efficient pump absorption despite the small absorption per double pass through the thin disk. A consequence of this is that the pump intensity is approximately constant along the beam axis.

In pulsed operation, the average intensity distribution in the thin disk is a standing-wave pattern with full contrast near the reflecting end. Because of the finite pulse bandwidth, the contrast decreases towards the other (anti-reflection-coated) end. At each location in the disk, the excitation level of the active ions is determined by the pump and local laser intensities. The spectral shape of the gain for a double pass through the disk is determined not only by the overall excitation in the gain medium, but is also directly affected by standing-wave effects because each spectral component of a beam is associated with a simple (full-contrast) standing-wave pattern. The period of this pattern depends on the frequency, and the local intensity of the mode pattern determines how strongly the spectral component interacts with the excited ions.

The dynamics of a laser with SHB in the gain medium can be very complicated because the lasing spectrum depends on the spectral shape of the gain, while the latter again reacts to changes in the pulse spectrum. We avoid these complications in our model by assuming a steady state with a given average laser power and soliton-pulse duration (and thus a given spectral width). In the experiment, the pulse duration and the shape of the spectrum are largely determined by the total (negative) intracavity dispersion and self-phase modulation, as we use soliton mode locking. For the given laser power and spectral width, we calculate the spectrum of the saturated gain and use this to determine whether or not the given state can be stable. This approach is numerically much more efficient than the simulation of the temporal evolution of arbitrary lasing spectra over a huge number of cavity round-trips (as used, for example, in [8]). Also we will show that it gives additional physical insight.

For the sake of the efficiency of our numerical calculations, we use a number of approximations. The thin disk is




Fig. 1. Geometry of a thin-disk laser head. The laser beam is reflected at the disk, and the pump optics arrange for, for example, eight double passes of the pump radiation through the disk

assumed to act as the end mirror of a standing-wave cavity, because a slight folding angle would make only a small difference in the period of the standing-wave pattern. Instead of the Gaussian transverse profile of the laser beams, we use a top-hat transverse intensity distribution so that the intensities are functions only of the coordinate z , the distance from the reflecting end of the thin disk. Effectively we average between the stronger effect of SHB on the beam axis and the weaker SHB in the wings of the spatial profile. The pump intensity is assumed to be constant, because the pump absorption in a single pass through the thin disk is small, and efficient pump absorption is achieved only through multiple passes of the pump radiation. Moreover, we approximate the gain spectrum of Yb:YAG near 1030 nm with a Gaussian function with a full width at half maximum (FWHM) of 1.55 THz or 5.5 nm. This quite accurately fits the gain spectrum near the gain maximum within the typical pulse bandwidth, whereas some deviations outside this bandwidth are not relevant. The laser transition is assumed to be homogeneously broadened. We also use a Gaussian function for the intracavity spectral power density $p_L(\nu)$, although this shape somewhat differs from the spectrum of a soliton (sech^2), which has stronger wings. In effect, we tend to slightly overestimate the effect of SHB. The Gaussian functions allow us to solve a substantial part of the problem analytically so that the numerical part becomes significantly less time-consuming. A final approximation is that we neglect energy migration in the Yb:YAG disk, which is difficult to quantify. Due to the relatively high Yb³⁺ doping level (typically 8–9 at.%), we expect some energy migration to occur, which should somewhat smoothen the distribution of excited ions and thus reduce the effect of SHB.

We now discuss the details of the calculations. The intracavity laser radiation is characterized by the spectral power density

$$p_L(\nu) = 2\sqrt{\frac{\ln 2}{\pi}} \frac{P_L}{\Delta\nu_L} \exp\left[-4 \ln 2 \left(\frac{\nu - \nu_0}{\Delta\nu_L}\right)^2\right], \quad (1)$$

a Gaussian function with the peak at frequency ν_0 and the FWHM $\Delta\nu_L$, which is normalized so that the intracavity laser power P_L is

$$P_L = \int_0^{+\infty} p_L(\nu) d\nu. \quad (2)$$

Because the gain per round-trip is moderate, we assume a constant power along the propagation direction. For each frequency component ν , there is a fully modulated standing-wave pattern, so that the spectral density of the local laser intensity (taking into account the superposition of the counter-propagating waves) is

$$i_L(z, \nu) = 2 \frac{p_L(\nu)}{A} \left[1 - \cos \frac{4\pi n \nu z}{c}\right]. \quad (3)$$

Here we have introduced the beam cross-section A and the refractive index n of the gain medium ($n = 1.82$ for YAG). We have assumed the phase jump at the reflecting coating to be π , i.e. to be independent of the frequency, which is reasonable

given the limited pulse bandwidth. The total laser intensity at a distance z from the reflecting coating is then

$$I_L(z) = \int_0^\infty i_L(z, \nu) d\nu, \quad (4)$$

and this integration can be performed analytically due to the Gaussian pulse spectrum. We obtain

$$I_L(z) = 2 \frac{P_L}{A} \left[1 - \exp \left(- \left(\frac{\pi n z \Delta \nu_L}{c \sqrt{\ln 2}} \right)^2 \right) \cos \frac{4\pi n \nu_0 z}{c} \right]. \quad (5)$$

This function represents a standing-wave pattern which is fully modulated for small z and continuously loses contrast for larger z .

Now we calculate the local excitation level, or more precisely the fraction N_2 of the Yb^{3+} ions which are in the excited state. (N_2 can vary between 0 and 1.) The laser cross-sections are assumed to depend on the frequency ν according to

$$\sigma_{\text{abs}}^{(L)}(\nu), \sigma_{\text{em}}^{(L)}(\nu) \propto \exp \left[-4 \ln 2 \left(\frac{\nu - \nu_0}{\Delta \nu_g} \right)^2 \right]. \quad (6)$$

We assume such a dependence for both the emission and absorption cross-sections, which is a good approximation within a bandwidth of a few nanometers, as can be seen from the McCumber relations [13]. The peak is at the same frequency ν_0 as the peak of the laser spectrum. In a first approximation, which is justified if the laser bandwidth is much smaller than the gain bandwidth, we would neglect the finite gain bandwidth and calculate the excitation level N_2 simply from the pump intensity and the total local laser intensity $I_L(z)$, using the laser cross-sections at the peak wavelength. A simple way to take into account the finite gain bandwidth is to replace $I_L(z)$ by an effective intensity $I_{L,\text{eff}}(z)$, calculated in the same way as $I_L(z)$ except that an additional Gaussian factor $\exp(-4 \ln 2 (\nu - \nu_0)^2 / \Delta \nu_g^2)$ is included in the integrand. This factor takes into account the reduced interaction for frequencies away from the center frequency. By combining the two Gaussians in the integral we obtain

$$I_{L,\text{eff}}(z) = 2 \frac{P_L}{A} \frac{\Delta \nu_{\text{eff}}}{\Delta \nu_L} \times \left[1 - \exp \left(- \left(\frac{\pi n z \Delta \nu_{\text{eff}}}{c \sqrt{\ln 2}} \right)^2 \right) \cos \frac{4\pi n \nu_0 z}{c} \right], \quad (7)$$

with the reduced bandwidth $\Delta \nu_{\text{eff}}^{-2} := \Delta \nu_L^{-2} + \Delta \nu_g^{-2}$. The normalized local excitation N_2 can then be calculated from the rate equations of a 3-level medium, using the effective intensity $I_{L,\text{eff}}(z)$ together with the peak cross-sections:

$$N_2(z) = \frac{\sigma_{\text{abs}}^{(P)} \tilde{I}_p(z) + \sigma_{\text{abs}}^{(L)} \tilde{I}_{L,\text{eff}}(z)}{\left(\sigma_{\text{abs}}^{(P)} + \sigma_{\text{em}}^{(P)} \right) \tilde{I}_p(z) + \left(\sigma_{\text{abs}}^{(L)} + \sigma_{\text{em}}^{(L)} \right) \tilde{I}_{L,\text{eff}}(z) + \tau_{\text{fl}}^{-1}}, \quad (8)$$

where the superscripts (P) and (L) refer to the pump and laser beams, respectively, the renormalized intensities are defined by $\tilde{I}^{(P,L)} = I^{(P,L)} / h\nu_{(P,L)}$, and all cross-sections apply to the frequency ν_0 . We have used the fact that all Yb^{3+} ions are either in the excited or in the ground-state Stark manifold. Note that (7) and (8) allow us to calculate without performing a time-consuming numerical integration over all frequencies.

From the excitation we obtain the frequency-dependent gain for a double pass through the disk:

$$G(\nu) = 2N_{\text{dop}} \int_0^d \left[N_2(z) \sigma_{\text{em}}^{(L)}(\nu) - (1 - N_2(z)) \sigma_{\text{abs}}^{(L)}(\nu) \right] \times \left[1 - \cos \frac{4\pi n \nu z}{c} \right] dz. \quad (9)$$

Here, N_{dop} is the doping density and d is the thickness of the disk. The oscillating factor in (9) accounts for the standing-wave pattern for the frequency ν , which modulates the interaction of a probe wave of frequency ν with the excited ions. This factor is responsible for the fact that the spectrum of the saturated gain deviates from the ‘natural’ gain spectrum of the Yb^{3+} ions. Finally, the effective gain for a pulse with spectral power density $p_L(\nu)$ is

$$G_{\text{eff}} = \frac{\int G(\nu) p_L(\nu) d\nu}{\int p_L(\nu) d\nu}. \quad (10)$$

This gain, calculated for the soliton pulse with given duration, must balance the cavity losses. In our model, we fulfill this condition by adjusting the pump intensity. We will later investigate in detail for which range of pulse bandwidths the resulting gain spectra $G(\nu)$ (from (9)) lead to stable situations, and how the stability range depends on various parameters.

We note that the integrand in (9) is a function with fast oscillations. For an efficient numerical evaluation we used the fact that the integrand is a nearly periodic function over a length of a few standing-wave periods. It is thus advisable to define a function which is the integrand averaged over one standing-wave period (by numerical integration), and then to integrate this relatively smooth function over the whole crystal length. In this way the saturated gain spectrum can be calculated within a few seconds on a standard personal computer, and a large number of configurations can be investigated within a reasonable time. Note that the algorithm used in [8] needs far more computation time, but of course it is more general in the sense that it can be applied to lasers which do not operate in the soliton mode-locked regime or are based on media with more complicated gain spectra.

2 Results for the Yb:YAG thin-disk laser

Here we quantitatively discuss a thin-disk Yb:YAG laser which has been slightly modified compared to the laser described in [1]. The setup is shown in Fig. 2. The output-coupler transmission has been increased to $T_{\text{out}} = 8\%$, and the total cavity loss is estimated to be 10% per round-trip. As the Yb:YAG disk (thickness $d = 220 \mu\text{m}$, refractive index $n = 1.82$) is a folding mirror in the standing-wave cavity, it amplifies a circulating pulse twice per round-trip, and the effective




Fig. 2. Setup of the thin-disk laser. DM = dispersive mirror, SESAM = semiconductor saturable absorber mirror. The prism near the output coupler is used to fine-tune the dispersion

gain per double pass through the disk must be 5%. The pump power is always adjusted to generate this gain. The mode area on the disk is $A = 10^{-6} \text{ m}^2$, corresponding to a mode radius of 0.56 mm. For an output power of $P_{\text{out}} = 15 \text{ W}$, we use an effective intracavity power of $2P_{\text{out}}/T_{\text{out}} = 375 \text{ W}$, where the factor 2 accounts for the fact that gain saturation occurs in two double passes per round-trip. Dispersion compensation is accomplished not with an adjustable (but slowly drifting) GTI (Gires–Tournois interferometer) as described in [1], but with 10 bounces per round-trip on three dispersive mirrors with a group-delay dispersion of $\approx -400 \text{ fs}^2$ at 1030 nm per bounce. An additional prism near the output coupler [14, 15] allows for the fine adjustment of the total intracavity dispersion and self-phase modulation and thus the soliton-pulse width. The SESAM is of a low-finesse design with 0.9% modulation depth, a saturation fluence of 0.1 mJ/cm^2 , and a recovery time of $\approx 100 \text{ ps}$ [1]. The pulse-repetition rate is 35 MHz.

The typical soliton-pulse duration of 0.63 ps corresponds to a bandwidth of 1.8 nm. In Fig. 3 we plotted the calculated spectra of the saturated gain per round-trip (two double passes through the disk) for different pulse bandwidths as well as the corresponding pulse spectra, and now discuss whether these cases correspond to stable situations. For the smallest bandwidth of 1 nm, a pronounced dip in the center of the gain spectrum occurs. This results from the fact that frequency components with an offset of $\approx 0.3 \text{ THz}$ ($\approx 1 \text{ nm}$) from the line center, still well within the natural gain bandwidth of 1.55 THz, can exploit regions with weak gain saturation (and thus high excitation level) near the AR-coated end of the crystal. This effect becomes weaker for broader pulse bandwidths. For the large pulse bandwidth of 2.5 nm, the dip disappears because the pulse spectrum is broad enough to largely ‘wipe out’ the standing-wave pattern near the AR-coated end. Still there is some broadening of the gain spectrum. Note that even in this situation there are still regions with weakly saturated gain near the reflecting end of the disk, but these could be probed only by frequency components outside the natural gain bandwidth. In Fig. 4 we plotted the minimum and maximum normalized excitation levels (within a standing-wave period) along the propagation direction for different pulse bandwidths. (The standing-wave period is too small to permit plotting the oscillating excitation level itself.) The excitation maxima are quite high near the reflecting end (left side of the graph), because in this region the standing-wave pattern is fully modulated and the gain is not saturated in the nodes of this pattern. For 2-nm pulse bandwidth the standing-wave pattern is largely wiped out near the AR-coated end of the




Fig. 3. Spectra of the saturated gain per round-trip (two double passes through the disk) for different pulse bandwidths, and the corresponding Gaussian pulse spectra (all dotted)




Fig. 4. Minimum and maximum normalized excitation levels along the propagation direction, for 1-nm (solid curves), 2-nm (dashed curves), and 3-nm (dotted curves) pulse bandwidth. The left and right parts of the graph correspond to the reflecting and the AR-coated ends of the crystal, respectively

crystal (right side of the graph). The pulse bandwidth required to wipe out the standing-wave pattern near the AR-coated end is inversely proportional to the disk thickness. It has been discussed, for example in [7], that this bandwidth is in the order of $c/(2nd)$, i.e. the free spectral range of the disk (if it were used as an etalon).

In the experiments, stable mode locking was achieved only for a pulse bandwidth in the range of approximately 1.7 nm to 2 nm. Outside this range, a rather unstable regime with a tendency for the generation of multiple pulses in the cavity (and large power fluctuations) was observed. In the following we will examine the reasons for this observation. Consider first the case with only 1-nm pulse bandwidth, which can not be stable for several reasons. First of all, a cw background with a frequency offset of $\approx 0.3 \text{ THz}$ will see 1.7% more gain per round-trip. This gain advantage of the cw background can not be compensated by a SESAM with only 0.9% modulation depth. Therefore we would expect that a cw background can grow in power and soon destabilize the soliton. Second, the soliton could split into two solitons each of half the energy, thus also half the bandwidth, and these two pulses could optimize their gain by slightly shifting their peak wavelengths. Third, the original soliton itself would not be spectrally stable because a shift of the spectrum to either side would increase the gain. Of course, any substantial change of the pulse spectrum will after some time result in changes of the gain

spectrum, and this is why the operation is unstable in such cases.

Now consider the case with a pulse bandwidth of 2.5 nm. As the dip in the saturated gain spectrum disappears, no pulse can optimize its gain by shifting its spectrum. However, a pulse with half the bandwidth, as obtained by splitting the soliton into two solitons with half the energy, experiences 0.43% more gain per round-trip. Such pulses have only slightly more loss at the SESAM due to the weaker absorber saturation. If the temporal distance of the pulses is smaller than the recovery time of the SESAM (≈ 100 ps), this difference of loss is even smaller, but only for the second pulse. Finally, a cw background has 0.5% more gain, but also experiences a substantially higher loss at the SESAM.

We have seen that several modes of operation can compete with the desired mode, which is a single-soliton pulse. The latter can be stable only if no other mode has a gain advantage which can not be compensated by the SESAM. We thus arrive at the following conditions for stable soliton mode locking:

1. A small shift of the pulse spectrum towards shorter or longer wavelengths should decrease the effective gain, so that the pulse is spectrally stable.
2. The difference between the maximum of the saturated gain curve (which could be exploited by a cw background) and the effective gain for the pulse must not exceed a certain value, which depends mainly on the modulation depth of the SESAM and its recovery time.
3. The difference between the effective gain for two solitons with half the bandwidth (and optimized spectral position) and the effective gain for the wanted soliton pulse must not exceed a certain value, which depends mainly on the modulation depth of the SESAM and the degree of SESAM saturation.

In the following, we discuss this in detail. In Fig. 5 we have plotted the gain advantage per round-trip (compared to the desired single pulse) for a cw background, for double pulses (with each pulse having half the energy, half the bandwidth, and optimized spectral position), and for a single pulse with optimized spectral position, all as functions of the pulse bandwidth. We see that according to condition 1, stable mode locking with less than 1.6-nm bandwidth is impossible, because the pulse would not be spectrally stable. We further see that a cw background (condition 2) always has the largest gain advantage, as to be expected, and that this curve has




Fig. 5. Gain advantage of cw background, double pulses, and pulse with optimized spectral position, all compared to the gain of a single-soliton pulse, as functions of the pulse bandwidth

a minimum for a pulse bandwidth of slightly below 2 nm. With a SESAM modulation depth of 0.9%, a gain advantage of the cw background in the order of 0.5% could be tolerated because much of the cw background experiences the loss of an essentially unsaturated SESAM. The gain advantage of double pulses (condition 3) is always smaller than for a cw background, and yet double pulses can more easily compete with the single pulse because they can saturate the SESAM nearly as much as the single pulse. For a quantitative estimate, we consider the energy loss $q(E_p)$ of a pulse on the SESAM (modulation depth ΔR and saturation energy E_{sat}) as a function of the pulse energy E_p [10]:

$$q(E_p) = \Delta R [1 - \exp(-E_p/E_{\text{sat}})] \frac{E_{\text{sat}}}{E_p}. \quad (11)$$

Note that this is not the loss of the absorber after passage of the pulse, but rather the fractional energy loss of the pulse itself. Equation (11) can be derived from the evolution equation

$$\frac{dq}{dt} = -\frac{P}{E_{\text{sat}}} q \quad (12)$$

for the loss of a slow saturable absorber, with $q(0) = \Delta R$ and $q \ll 1$. Introducing the saturation parameter $S = E_p/E_{\text{sat}}$, we obtain

$$q(S) = \Delta R (1 - \exp(-S)) / S, \quad (13)$$

which is roughly $\Delta R/S$ for $S \geq 3$. If the SESAM is operated with a typical value of, for example, $S = 8$, the loss of a pulse with half the energy (i.e. $S = 4$), will be larger by $\approx \Delta R/8$, which is $\approx 0.1\%$ in our case. From Fig. 5 we see that only for a pulse bandwidth near 1.8 nm, this loss difference is large enough to suppress the formation of double pulses. This is in remarkable agreement with the mentioned experimental observations.

Note that even a different type of soliton mode-locked laser without SHB would exhibit the described instability according to condition 3 at large pulse bandwidths, and this instability usually sets a lower limit to the achievable pulse duration. For a laser with weak SESAM saturation, condition 2 may be the limiting factor. The instability for small pulse bandwidths (long pulses), however, occurs only with SHB and again usually acts through condition 3. It is instructive to examine Fig. 6, which shows the same curves as Fig. 5 but




Fig. 6. Similar to Fig. 5, but without the effect of spatial hole burning: gain advantage of cw background and double pulses, both compared to the gain of a single-soliton pulse, as functions of the pulse bandwidth

for a hypothetical laser with the same parameters but without SHB. Here, the gain advantage of both a cw background and double pulses becomes very small for a small enough pulse bandwidth. On the other hand, for a large pulse bandwidth (e.g. 2 nm) such a laser would be more difficult to mode-lock than the laser with SHB. With the given SESAM, we would expect to achieve at most a 0.9-nm pulse bandwidth, or 1.2-ps soliton-pulse duration. This is because SHB leads to a saturated gain spectrum which is broader than the natural gain spectrum. As this spectrum still has a quite smooth shape, stable bandwidth-limited solitons can be formed.

Returning to the laser with SHB, we have shown that stable mode locking can be achieved only in a certain range of pulse bandwidths and thus pulse durations, with the width of this range depending on the parameters of the gain medium (cross-sections, gain bandwidth, and crystal length), the intracavity intensity, and the SESAM parameters. Apparently there could be situations where stable mode locking is not possible at all, and the described Yb:YAG thin-disk laser is indeed not far from such a point. It is therefore interesting to know how the stability range could be expanded. An obvious choice is the use of a SESAM with larger modulation depth, which is more capable of discriminating against cw background and double pulses. However, any significant increase beyond the currently used 0.9% would lead to Q-switched mode locking (QML) [10,11]. The same problem excludes the possibility of using a larger spot size on the SESAM, which would increase E_{sat} , decrease the saturation parameter S , and thus help to suppress double pulses. A somewhat faster recovery time would help to suppress a growing cw background, but not double pulsing, which seems to be the more severe problem. Only a recovery time below the pulse duration (i.e. a fast saturable absorber) would help, because this would strongly favor a single pulse (against double pulses) due to its higher peak intensity. In a semiconductor absorber, a certain part of the saturation recovers by intraband thermalization within ≈ 100 fs, but another part results from carrier recombination and trapping, which can not occur on a sub-picosecond time scale, except if the material is grown at rather low temperatures and increased non-saturable losses are accepted [16,17]. We therefore see that only a fast absorber, based, for example, on the Kerr effect, could in principle allow for mode locking in a broader range of pulse durations, although Kerr-lens mode locking introduces additional difficulties, in particular operation of the laser cavity near a region of instability. On the other hand, no slow saturable absorber could be expected to be significantly superior to the SESAM used in our experiments, because the crucial parameters (modulation depth and saturation parameter) are already optimized.

We also investigated the effect of reducing the total cavity loss from 10% to 5%, where we assumed the average output power to stay unchanged. With half the gain, gain-filtering effects are reduced. On the other hand, gain saturation in the thin disk and thus the SHB effect become stronger. Figure 7 shows that the latter effect is less important: the gain advantages for both double pulses and a cw background are significantly reduced. However, in order to take advantage of this, the spot size on the SESAM should be increased because otherwise the saturation parameter S would rise (as an effect of the increased intracavity pulse energy) and the discrimination against double pulses would be reduced. Indeed we observed

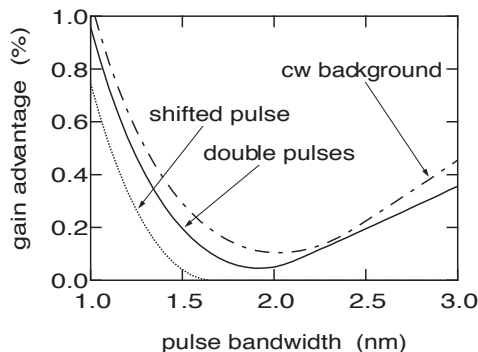


Fig. 7. Similar to Fig. 5, but for reduced cavity losses of 5% (instead of 10%)

in experiments that modifying the output-coupler transmission alone has little effect on the stability range for mode locking. Also note that a decreased output-coupler transmission helps to avoid Q-switching instabilities [11] and thus would allow us to employ a SESAM with higher modulation depth, which would significantly expand the stability range. However, this also tends to reduce the laser efficiency if parasitic cavity losses can not be reduced.

We also consider modifying the crystal thickness. The further development of high-power diode lasers with better spatial beam quality will allow us to construct thin-disk laser heads with an increased number of passes of the pump radiation through the disk. This allows for a further reduction of the disk thickness, with beneficial effects on the thermal properties, mainly a further reduction of thermal lensing. Figure 8 shows how various parameters change when the crystal thickness is reduced. The effective pump power (adding the contributions from all passes) is always adjusted to achieve the same gain of 5% per double pass through the disk, and is indicated by the circles. The squares show the pulse bandwidth for which the gain advantage for double pulses (usually the main challenge) is at its minimum, while the triangles indicate how large this gain advantage is at this point. Compared to the original situation with 220- μm disk thickness, a reduced thickness of ≈ 100 μm would allow for slightly shorter pulses, while the stability against double pulsing would be slightly reduced. However, for a further reduced disk thickness of, for example, 50 μm , the stability against double puls-

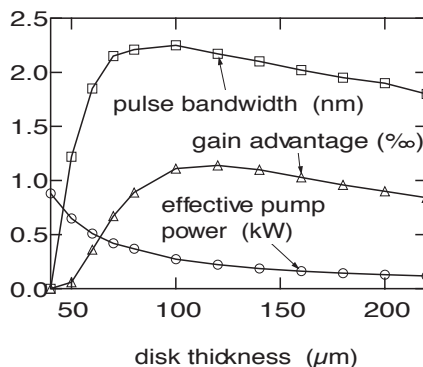


Fig. 8. Variation of the disk thickness. Rectangles: pulse bandwidth with minimum gain advantage for double pulses. Triangles: minimum gain advantage (in $\%_0 = 0.1\%$) for double pulses. Circles: effective pump power needed to generate 5% of gain per double pass

ing becomes much better, while the optimum pulse duration would then be longer (≈ 1 ps). With 40- μm disk thickness or less, arbitrarily long pulses become possible. In this regime, which may be reached in the future, the effect of SHB becomes weaker because there are no more wavelength components within the gain bandwidth which can exploit the regions of undepleted gain.

A higher doping level of the crystal would also allow us to reduce the crystal thickness. In addition, it should increase the effect of energy migration and thus reduce the effect of SHB. On the other hand, it can reduce the laser efficiency by enhancing quenching effects. The currently used doping level (9%) already seems to be close to the optimum.

Another parameter which can be modified is the cavity length, which determines the repetition rate of the laser. If we make the cavity longer, this does not directly affect the SHB effect which depends only on the average intensity in the thin disk. However, the intracavity pulse energy rises (if the output-coupler transmission is not changed), and the spot size on the SESAM should be increased to keep the saturation parameter S unchanged. Still the Q-switching tendency is reduced, as can be shown with the equations of [11], so that a SESAM with higher modulation depth can be employed. The thereby increased heat load on the SESAM is no problem as the spot size is also increased; it only somewhat decreases the laser efficiency. The higher modulation depth finally gives more stability against double pulsing. Thus we find that the double-pulsing instability can be reduced by increasing the modulation depth of the SESAM and using a relatively long laser cavity. Indeed our Yb:YAG thin-disk laser has already been designed to operate with a quite low repetition rate of 35 MHz.

3 Results with other gain media

We finally discuss the effects of spatial hole burning on the mode locking of thin-disk lasers with other gain media, having larger or smaller amplification bandwidths. We will see that the effect of SHB is very different in these cases and mainly depends on the parameter $\gamma := nd\Delta\nu_g/c$, which is half the ratio of the FWHM gain bandwidth and the free spectral range of the disk (if used as an etalon). For the Yb:YAG laser, we have $\gamma = 2.1$. Here, the natural gain bandwidth is just sufficient to form pulses which can largely wipe out the standing-wave pattern in the disk. For Nd:YAG, where the gain bandwidth is only 0.5 nm, we get a much smaller value of γ . For the modeling we use numbers which are similar to those in previous experiments with continuous-wave TEM₀₀ thin-disk Nd:YAG lasers [18]. (A passively mode-locked thin-disk Nd:YAG laser has not yet been demonstrated.) Because of the weaker pump absorption, a larger disk thickness of 350 μm is needed, which leads to $\gamma = 0.28$. The laser is assumed to have 12 W of average output power with an output-coupler transmission of 5%. The laser-beam radius in the thin disk is 1.1 mm. As in the previous cases, we plotted the gain advantage of a cw background and of double pulses as a function of the spectral width of the laser emission (Fig. 9). As in the hypothetical case of Yb:YAG without SHB (Fig. 6), these gain advantages become very small as the pulse bandwidth gets small, so that stable mode locking could be achieved for arbitrarily long pulses. Figure 10 shows that




Fig. 9. Similar to Fig. 5, but for a Nd:YAG laser




Fig. 10. Spectrum of the saturated gain per round-trip for a small pulse bandwidth of 0.05 nm in a Nd:YAG laser (solid curve) and, for comparison, the natural Nd:YAG gain spectrum (dashed curve). Dotted curve: pulse spectrum

a dip in the saturated gain spectrum can not be formed, even for very narrow pulse spectra, because for all wavelengths within the gain bandwidth, the standing-wave patterns are so similar that none of these waves can exploit the regions of undepleted gain in the crystal. There is only some broadening of the gain spectrum. The lasers considered in [7, 8] all operated in this regime.

The situation is different for Nd:YVO₄. This medium has about twice the gain bandwidth of Nd:YAG. Due to the large cross-sections, a somewhat thinner disk can be used. We simulated a case with a 280- μm thin disk ($\gamma = 0.52$), a beam radius of 1.1 mm in the disk, and 10 W of output power through an output coupler with 20% transmission (increased compared to the Nd:YAG laser because of the higher achievable gain). The results (Fig. 11) show that stable mode locking should be possible for a pulse bandwidth near 0.32 nm, corresponding to a pulse duration of about 3.7 ps. The minimum required modulation depth of the SESAM is $\approx 2\%$, somewhat higher than for the discussed Yb:YAG laser. However, this should not constitute a problem since the intracavity power is relatively small and the Q-switching tendency is weak due to the large cross-sections of Nd:YVO₄. For this reason, we expect the mode locking of a thin-disk laser with Nd:YVO₄ to be easier to achieve than with Yb:YAG. We also note that with a reduced disk thickness one could generate arbitrarily long pulses, as discussed for Yb:YAG in Sect. 2.

We also briefly discuss the possibility of mode locking thin-disk lasers based on gain media with a larger gain bandwidth compared to Yb:YAG. This might lead to high-power

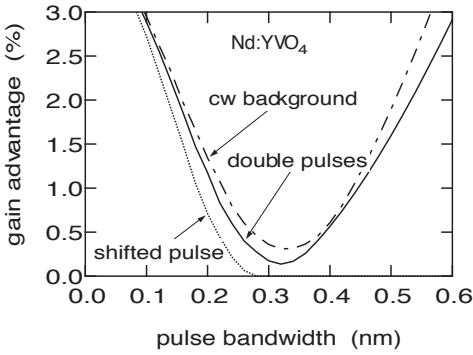


Fig. 11. Similar to Fig. 5, but for a Nd:YVO₄ laser

lasers with pulse durations of 200 fs or below. The main challenge is that such gain media typically have relatively poor thermal properties. Also, large enough laser cross-sections are required to suppress QML. In addition, large enough pump and laser cross-sections and a high enough doping level are needed in order to keep the disk thickness reasonably small. Finally, the medium has to have sufficient mechanical strength so that a thin disk can be mounted on the cooling finger. At this point it is not yet clear which material will be most suitable for a < 0.5-ps thin-disk laser, although Yb-doped tungstate crystals like Yb:KGW or Yb:KYW [19] seem to be the most promising candidates, with a small quantum defect, even slightly higher laser cross-sections than Yb:YAG, and a fairly large bandwidth. Also, some other parameters (disk thickness, mode sizes, output coupling etc.) will have to be optimized for such a laser and are not yet known. We thus investigate the effect of SHB in a laser based on a hypothetical gain medium, which is assumed to have the same properties as Yb:YAG except that the gain bandwidth is 40 nm (instead of 5.5 nm). We also assume the same operating parameters as for the first Yb:YAG laser discussed in Sect. 2 and obtain $\gamma = 17$. Figure 12 shows that stable mode locking should be easily achieved for ≈ 7 -nm pulse bandwidth (corresponding to ≈ 160 -fs pulse duration). Here the gain advantage particularly for double pulses is significantly weaker than for the original Yb:YAG laser. This is because the standing-wave pattern in the disk can easily be wiped out with a bandwidth which is still smaller than the natural gain bandwidth. On the other hand, SHB does not significantly reduce the possible pulse duration in this regime: the improvement compared to Yb:YAG is smaller




Fig. 12. Similar to Fig. 5, but for a hypothetical laser material like Yb:YAG except that the gain bandwidth is 40 nm

than may be naively expected from the increase of amplification bandwidth.

4 Methods to reduce the effect of spatial hole burning

We have seen that spatial hole burning in thin-disk lasers can help to generate shorter pulses, but also can strongly limit the range of pulse durations which can be realized. It would therefore be desirable to eliminate this effect in some cases, for example if the aim is the generation of longer pulses, or if a larger range of pulse durations is required.

It is clear that using a ring cavity will not have the desired effect as in a thin-disk laser there is always the superposition of counter-propagating waves in the gain medium. One might also think of employing polarization effects. Any polarization state can be decomposed into two orthogonal linear components, and these can be seen as generating two independent standing-wave patterns in the disk. The phase of these standing-wave patterns is determined by the boundary conditions at the reflecting surface of the disk, so that both standing waves are always in phase (and thus can not cancel each other) unless a polarization-dependent phase change could be introduced near the reflecting end. For this purpose one could bond the thin disk to a quarter-wave plate and make the reflecting coating on the latter [20]. As the heat then has to flow through the quarter-wave plate, thermal lensing effects are increased. This effect could be limited by making a thin zero-order wave plate from a strongly birefringent material with good thermal conductivity.

In Sects. 4.1 and 4.2, we propose two other methods and investigate their potentials.

4.1 Use of a filter

As the problem – unstable mode-locking behavior outside a narrow range of pulse durations – results from a distortion of the saturated gain spectrum, it appears natural to counteract this by using an additional wavelength filter in the laser cavity. It is obvious that this method should work if the filter response is designed so as to flatten the effective gain spectrum. However, it is not obvious to what extent this method can work with filters of a given simple spectral shape, how tight the fabrication tolerances would be, and how large the range of possible pulse duration becomes. Therefore we studied various cases by including different filter functions in the model described above.

4.1.1 Gaussian filter. One of the simplest filter functions, which is often approximated by real filters, is a Gaussian where the filter transmission is

$$T_F(\nu) = \exp \left[-4 \ln 2 \left(\frac{\nu - \nu_0}{\Delta \nu_F} \right)^2 \right]. \quad (14)$$

The only variable parameter is the FWHM filter bandwidth $\Delta \nu_F$. We first investigated the behavior for an Yb:YAG laser with the same parameters as the first laser modeled in Sect. 2. In Fig. 13, we plotted the gain advantage for double pulses versus pulse bandwidth for various values of $\Delta \nu_F$. All curves have been truncated for short wavelengths, where the pulses




Fig. 13. Similar to Fig. 5, but with additional Gaussian filters with FWHM filter bandwidths of 5.8 nm, 10 nm, 17 nm, 25 nm, 33 nm, 50 nm, and without a filter (from left to right)

would not be spectrally stable. First considering relatively weak filters (bandwidth of 17 nm or larger), we see that the optimum pulse bandwidth is reduced with decreasing filter bandwidth, as one would expect. However, the minimum gain advantage for double pulses is not reduced, but even slightly increased. This means that the use of a Gaussian filter would allow us to generate longer pulses, but an absorber with a slightly increased modulation depth would be required, and the stability range would be further reduced. Certainly the filter will not make stable mode locking easier to achieve. Also note that for a filter with, for example, 17-nm bandwidth the filter loss would make it necessary to increase the pump power by 16% to achieve the same output power as without the filter. This is partly because of the loss at the filter and partly because for narrow-band emission there are regions of largely undepleted gain in the disk.

Note, however, that for considerably smaller filter bandwidths the situation becomes totally different. Then the situation is similar to the Nd:YAG laser of Sect. 3: there are no more frequencies within the filter bandwidth which have a sufficiently modified standing-wave period to make use of regions with undepleted gain. For a filter bandwidth of ≈ 5.8 nm, stable mode locking should be possible for any pulse bandwidth between 0 and 0.24 nm (Fig. 13). (Filters with less than 5.8-nm bandwidth would further reduce the achievable pulse bandwidth.) The filter is now strong enough to prevent the formation of a dip in the saturated gain spectrum even for an arbitrarily small pulse bandwidth. (Note that the spectral width of the dip is mainly determined by the free spectral range of the disk, not by the pulse bandwidth.) Interestingly, the filter will not introduce any significant loss if the pulse bandwidth is below ≈ 0.05 nm, and yet serves to stabilize the pulses against the growth of other spectral features. Nevertheless the laser efficiency is significantly reduced because a narrow-band beam can not access all the excited ions in the disk. The pump power has to be increased by $\approx 9\%$ compared to the case without the filter and 1.8-nm pulse bandwidth, for example, in order to obtain the same average output power.

4.1.2 Fabry–Perot filter. The reason why a Gaussian filter (as discussed in Sect. 4.1.1) with moderate filter bandwidth does not provide more stability against double pulsing, can be understood as follows. In order to remove the dip in the saturated gain spectrum, the filter bandwidth can not be much

larger than the free spectral range of the disk. For such a narrow bandwidth, however, the filter inevitably increases the steep drop of gain outside the two maxima of the spectrum. Consequently, the double-pulsing instability limits the pulse bandwidth to quite small values.

For more stability against double pulsing in a regime with larger bandwidth, a smoother shape of the net-gain spectrum is required. To achieve this, we need a filter which helps to attenuate the gain maxima near the dip without increasing the drop of gain outside this region. It is apparent that a Fabry–Perot filter could fulfill this requirement if its free spectral range is chosen appropriately (somewhat larger than that of the disk) and a resonance occurs for the center frequency of the pulses. We have also modeled this case and found that for optimized filter parameters this indeed works very well. For example, we set the Fabry–Perot mirror reflectivities to be 0.36% and its free spectral range to be 2.2 nm. Here, a pulse bandwidth of 1 nm is achieved, and the gain advantage for double pulses is reduced to 0.05%, which is significantly less than without the filter (0.08% in the optimum case of 1.8-nm bandwidth, or 1.5% for 1-nm bandwidth). The resulting net-gain spectrum (Fig. 14) is quite flat. The filter loss is acceptable; it raises the demand of pump power for the same output power by 7%. Note, however, that the fabrication tolerances for the Fabry–Perot filter would be quite tight: the reflectivity should be $0.36 \pm 0.05\%$.

We also tried to optimize the parameters of the Fabry–Perot filter for a smaller pulse bandwidth of 0.2 nm and found that this does not work. The reason is that for this narrow bandwidth the shape of the spectral response of a Fabry–Perot filter is no longer suitable to produce an approximately flat net-gain spectrum.

4.2 Use of different reflection angles

Here we propose another scheme which significantly reduces the effect of SHB, without requiring additional specially designed optical components. The principle is to use a specially designed laser cavity with multiple bounces of the laser beam on the disk, with the reflection angles chosen so that the standing-wave pattern is at least partially wiped out. It turns out that for a typical disk thickness of 220 μm a difference in angles of a few degrees, which can easily be achieved, is sufficient for this purpose. Figure 15 shows an example




Fig. 14. Net-gain spectrum (*solid curve*) of an Yb:YAG laser for a pulse bandwidth of 1 nm and an optimized Fabry–Perot filter. Also shown is the natural gain spectrum (*dashed curve*) and the pulse spectrum (*dotted curve*)




Fig. 15. Example for a cavity design with two different reflection angles on the disk. M1 and M2 are two flat mirrors, which can be the output coupler and the SESAM. Dispersive mirrors can be used to adjust the cavity dispersion

for the design of such a cavity. It is interesting to note that this method alone would not be sufficient to achieve single-frequency operation of a thin-disk laser, because in single-frequency operation there is always a standing-wave pattern in the disk, which only becomes more complicated if a cavity with several reflections on the disk is used. In a mode-locked laser, however, the cavity arm lengths are so much longer than the spatial extent of the pulses that there can be no interference of the contributions from the passes with different reflection angles. We can thus calculate the pattern of excitation in the disk just by adding the intensity (not amplitude) patterns which result from the passes with different angles. This fact, which is easy to understand in the time domain, could also be explained in the frequency domain, although only in a more sophisticated way, which we do not discuss here.

A rough estimate of the required reflection angles can be obtained as follows. Each beam with a reflection angle α_j (in the crystal, relative to the normal direction) and a vacuum wavelength λ creates a standing-wave pattern in the disk with a period

$$\Lambda_j = \frac{\lambda}{2n \cos \alpha_j}. \quad (15)$$

If we choose a cavity with two angles α_1 and α_2 so that

$$\left| \frac{d}{\Lambda_1} - \frac{d}{\Lambda_2} \right| = \frac{2nd}{\lambda} |\cos \alpha_1 - \cos \alpha_2| = 1/2, \quad (16)$$

the corresponding standing-wave patterns will be just out of phase near the non-reflecting end of the disk (with thickness d). For a 220- μm -thick disk this condition is fulfilled, for example if $\alpha_1 = 1^\circ$ and $\alpha_2 = 2.29^\circ$. Then the resulting excitation pattern is smooth near the AR-coated end, if the powers in both beams are the same. Still there remains a significant oscillation of the excitation level in the middle of the crystal. It is thus better to choose a somewhat larger difference of angles.

For a more comprehensive study we generalize the model described in Sect. 1. We consider two beams with different angles, where the total intensity pattern is calculated by adding the contributions for both angles, as discussed above. The resulting gain is also the sum of the contributions from the two angles. We then have chosen a parameter set for a Yb:YAG laser (the most critical case) which is similar to the first parameter set discussed in Sect. 2, except that we assume a standing-wave cavity where the pulse passes the disk four times per round-trip, with two different angles α_1 and α_2 . The cavity loss is doubled, i.e. 20%, because we have twice the small-signal gain per round-trip and can thus choose a higher output-coupler transmission. The modulation depth




Fig. 16. Gain advantage of cw background versus the external angle β for fixed $\beta_1 = 2^\circ$. (The gain advantage for double pulses is very similar.) All curves have been truncated for short wavelengths, where the pulses would not be spectrally stable

of the SESAM should also be doubled to 2%. We keep α_1 fixed at 1.1° (corresponding to $\beta_1 = 2^\circ$ outside the crystal) and vary the external angle β_2 of the second beam.

First we consider a beam with narrow bandwidth (0.1 nm), corresponding to relatively long pulses (11 ps). Figure 16 shows the gain advantage of a cw background (which is hardly different from the gain advantage for double pulses in this case) as a function of the angle β_2 . While (16) would suggest a value of $\beta_2 = 4.24^\circ$, we see that the optimum is at $\beta_2 = 4.8^\circ$. Here, the gain advantage of the cw background is 1.2%, which is much less than with $\beta_1 = \beta_2$. However, the gain advantage of double pulses is also about 1.2%, and this is roughly six times more than can be tolerated when the modulation depth of the SESAM is in the order of 2%. A significantly higher modulation depth can not be employed because of Q-switching instabilities. We thus see that the use of two different angles is still not enough to allow for mode locking with arbitrarily long pulses.

The situation gets less critical if the pulse bandwidth becomes larger. In Fig. 17 we have plotted the gain advantage for double pulses versus pulse bandwidth for different values of β_2 . As in Fig. 13, all curves have been truncated for small bandwidths, where the pulses would not be spectrally stable. We see that the best situation is achieved for $\beta_2 = 5^\circ$ and a pulse bandwidth around 1.1 nm. Here the gain advantage for double pulses is only $\approx 0.034\%$, i.e. much less than without the use of two different reflection angles. With 2% modulation depth, about 0.2% gain advantage for double pulses




Fig. 17. Gain advantage of double pulses versus pulse bandwidth for fixed $\beta_1 = 2^\circ$ and various angles β_2

could be tolerated. This means that, for example, for $\beta_2 = 5^\circ$ one can achieve stable mode locking with a pulse bandwidth in the range of 0.95 nm to 1.47 nm, corresponding to pulse durations in the range of ≈ 0.8 ps to 1.2 ps. This is a significantly improved stability range compared to the situation with only one reflection angle.

For the generation of even longer pulses, this method can be combined with the use of a Gaussian filter as discussed in Sect. 4.1.1. For example, Fig. 18 shows the gain advantages of double pulses and a cw background for the same situation as in Fig. 17, but with an additional Gaussian filter of 20-nm bandwidth. With a SESAM of 2% modulation depth, a gain advantage for double pulses of about 0.2% should be tolerable, so that stable mode locking should now be possible for pulse bandwidths between 0.4 nm and 0.75 nm, corresponding to pulse durations between 1.5 ps and 2.8 ps. For less than 0.4-nm bandwidth, the pulses would not be spectrally stable.

If we further reduce the filter bandwidth to 14 nm (Fig. 19), the stability range gets much larger: now the pulse bandwidth can be varied anywhere between 0 and 0.54 nm if the condition is that the gain advantage for double pulses should be smaller than 0.2%. We compare this to the situation of Fig. 13 (Sect. 4.1.1). There, a filter with < 6 -nm bandwidth was required, and the maximum pulse bandwidth was only 0.24 nm. Also, the reduction of laser efficiency is significantly weaker if we use two reflection angles, because in this case we can access most of the excited ions in the disk.

Of course, another possibility would be to use a cavity with three (or even more) different reflection angles on the disk. This should lead to further increased stability ranges of

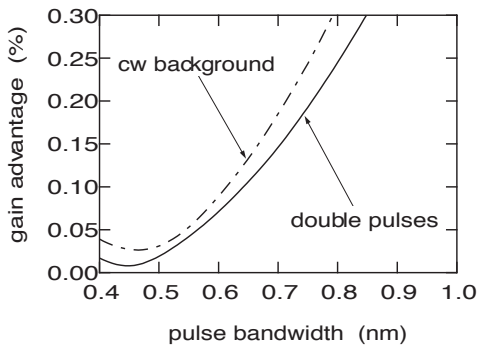


Fig. 18. Gain advantage of double pulses (*solid curve*) and cw background (*dashed curve*) versus pulse bandwidth for an Yb:YAG laser with two reflection angles ($\beta_1 = 2^\circ$, $\beta_2 = 5^\circ$) and a Fabry–Perot filter with 20-nm bandwidth




Fig. 19. Same as Fig. 18, but for a filter bandwidth of 14 nm

mode locking. However, this approach is more complex, and we do not discuss it further here.

5 Conclusions

We have investigated the effect of spatial hole burning on the mode-locking behavior of thin-disk lasers based on different gain materials. We have shown that this depends very much on the parameter $\gamma := nd\Delta\nu_g/c$, which is half the ratio of the FWHM gain bandwidth and the free spectral range of the disk (if used as an etalon). For the Yb:YAG laser which has already been experimentally demonstrated, we have $\gamma = 2.1$, where the pulse bandwidth can be just sufficient to largely wipe out the standing-wave pattern. The calculations explain with good quantitative agreement why stable mode locking can be achieved only in a narrow range of pulse durations. The generation of longer pulses, i.e. with smaller bandwidth, is prohibited by the formation of a dip in the inhomogeneously saturated gain spectrum, which can favor both a cw background and double pulses as competitors to the desired single-soliton pulse. On the other hand, shorter pulses are prohibited by the limited gain bandwidth, much as is the case in any other soliton mode-locked laser. The minimum achievable pulse duration is shorter than it would be without spatial hole burning. In the case of Yb:YAG, this improvement is a factor of the order of 2. The simulations as well as experimental results [1] show that in the soliton mode-locked regime this reduction of pulse duration is possible without a deterioration of the time–bandwidth product.

For Nd-doped gain media, typical values of γ are well below 1. For Nd:YAG, with $\gamma = 0.28$, a dip in the saturated gain spectrum can not be formed, even for narrow pulse spectra, and stable mode locking should be achievable for arbitrarily long pulses. The lasers discussed in [7,8] all operated in this regime, which may also be reached with Yb:YAG if the disk thickness is reduced to $\approx 50 \mu\text{m}$. For Nd:YVO₄, with a larger bandwidth than Nd:YAG and a somewhat larger value of $\gamma = 0.52$, SHB is important and again allows only for pulse durations in a certain range, in this case around 3.7 ps. Compared to Yb:YAG, the stability requirement can be more easily achieved because the much higher laser cross-sections permit the use of a saturable absorber with higher modulation depth, without Q-switching instabilities occurring.

On the other hand, Yb-doped gain media with larger gain bandwidths (and thus much larger values of γ) should allow stable mode locking with < 200 -fs pulse duration, provided that a medium with sufficiently good thermal and mechanical properties and also with large enough laser cross-sections is found. For such a material, the tendency for double pulsing is expected to be significantly weaker than for Yb:YAG, provided that the optimum pulse bandwidth is chosen. Longer pulses, however, can again not be generated in such configurations.

We have also discussed various methods to reduce the effect of SHB. This is of interest particularly if the generation of longer pulses or a larger range of possible pulse durations is desired. We found that in the case of Yb:YAG a Gaussian filter in the laser cavity will allow for the generation of longer pulses, but will not increase the range of possible pulse durations or increase the stability against double pulsing, except if the filter bandwidth is below a certain value

where arbitrarily narrow pulse bandwidths are possible. This regime is similar to the one for gain media with small amplification bandwidth (i.e. small values of the parameter γ). Better results even with relatively short pulses could be obtained with optimized Fabry–Perot filters where, however, the fabrication tolerances would be rather tight. Finally, we proposed to use laser cavities where different reflection angles on the disk occur. This significantly expands the range of possible pulse bandwidths, although arbitrarily long pulses can still not be generated in the considered cases. For arbitrarily long pulses, a filter can be used in addition, without the filter parameters being very critical. Compared to the situation with the filter alone, the stability range of mode locking as well as the laser efficiency are then significantly improved.

Acknowledgements. We thank T. Südmeyer for assistance with some of the measurements.

References

1. J. Aus der Au, G.J. Spühler, T. Südmeyer, R. Paschotta, R. Hövel, M. Moser, S. Erhard, M. Karszewski, A. Giesen, U. Keller: *Opt. Lett.* **25**, 859 (2000)
2. A. Giesen, H. Hügel, A. Voss, K. Wittig, U. Brauch, H. Opower: *Appl. Phys. B* **58**, 363 (1994)
3. M. Karszewski, U. Brauch, K. Contag, S. Erhard, A. Giesen, I. Johannsen, C. Stewen, A. Voss: In *Advanced Solid State Lasers*, Vol. XIX of OSA Technical Digest Series (Optical Society of America, Washington, DC 1998) p. 296
4. U. Keller, D.A.B. Miller, G.D. Boyd, T.H. Chiu, J.F. Ferguson, M.T. Asom: *Opt. Lett.* **17**, 505 (1992)
5. U. Keller, K.J. Weingarten, F.X. Kärtner, D. Kopf, B. Braun, I.D. Jung, R. Fluck, C. Hönninger, N. Matuschek, J. Aus der Au: *IEEE J. Sel. Top. Quantum Electron.* **2**, 435 (1996)
6. U. Keller: In *Nonlinear Optics in Semiconductors*, ed. by E. Garmire, A. Kost (Academic, Boston 1999) Chap. 4, 59, 211
7. B. Braun, K.J. Weingarten, F.X. Kärtner, U. Keller: *Appl. Phys. B* **61**, 429 (1995)
8. F.X. Kärtner, B. Braun, U. Keller: *Appl. Phys. B* **61**, 569 (1995)
9. F.X. Kärtner, U. Keller: *Opt. Lett.* **20**, 16 (1995)
10. F.X. Kärtner, L.R. Brovelli, D. Kopf, M. Kamp, I. Calasso, U. Keller: *Opt. Eng.* **34**, 2024 (1995)
11. C. Hönninger, R. Paschotta, F. Morier-Genoud, M. Moser, U. Keller: *J. Opt. Soc. Am. B* **16**, 46 (1999)
12. S. Erhard, K. Contag, I. Johannsen, M. Karszewski, T. Rupp, C. Stewen, A. Giesen: In *Advanced Solid State Lasers*, Vol. XXVI of OSA Technical Digest Series (Optical Society of America, Washington, DC 1999) p. 38
13. D.E. McCumber: *Phys. Rev.* **136**, 954 (1964)
14. M. Ramaswamy-Paye, J.G. Fujimoto: *Opt. Lett.* **19**, 1756 (1994)
15. D. Kopf, G.J. Spühler, K.J. Weingarten, U. Keller: *Appl. Opt.* **35**, 912 (1996)
16. M. Haiml, U. Siegner, F. Morier-Genoud, U. Keller, M. Luysberg, P. Specht, E.R. Weber: In *Symposium on Non-Stoichiometric III-V Compounds*, ed. by P. Kiesel, S. Malzer, T. Marek (Lehrstuhl für Mikrocharakterisierung, Friedrich-Alexander Universität, Erlangen-Nürnberg 1998) 6, 79
17. M. Haiml, U. Siegner, F. Morier-Genoud, U. Keller, M. Luysberg, R.C. Lutz, P. Specht, E.R. Weber: *Appl. Phys. Lett.* **74**, 3134 (1999)
18. A. Giesen, G. Hollemann, I. Johannsen: 1999 OSA Technical Digest Series (Optical Society of America, Washington, DC 1999) p. 29
19. N.V. Kuleshov, A.A. Lagatsky, V.G. Shcherbitsky, V.P. Mikhailov, E. Heumann, T. Jensen, A. Diening, G. Huber: *Appl. Phys. B* **64**, 409 (1997)
20. E. Sorokin: private communication (2000)