Optimization of chirped mirrors

Vladislav Yakovlev and Gabriel Tempea

We demonstrate that a highly efficient global optimization of chirped mirrors can be performed with the memetic algorithm. The inherently high sensitivity of chirped-mirror characteristics to manufacturing errors can be reduced significantly by means of the stochastic quasi-gradient algorithm. The applicability of these algorithms is not limited to chirped mirrors. © 2002 Optical Society of America
OCIS codes: 310.6860, 230.4170, 320.0320.

1. Introduction

Dispersive dielectric multilayer mirrors[1] [henceforth called chirped mirrors (CMs)] have contributed significantly to enhancement of the performance, compactness, and reliability of femtosecond laser sources. The generation of sub-10-fs pulses directly from oscillators requires accurate, broadband, higher-order dispersion control,[2] which can now be achieved routinely with CM dispersion-controlled oscillators.[3] Progress in CM design and manufacturing in conjunction with advanced oscillator architectures has permitted the direct generation of sub-6-fs pulses.[4–6] Novel spectral broadening techniques[7–9] allow spectra to be produced that can extend over more than 1 optical octave; subsequent compression with CMs has led to pulses shorter than 5 fs.[7,10,11,12]

The first CM designs were obtained by computer optimization and had a bandwidth of 200 nm at the central wavelength of 800 nm.[1] Several predesign methods have been proposed for the calculation of starting structures that, after limited computer optimization, converge to designs with enhanced performance.[13–16] Although the proposed design methods are substantially different from one another, the final designs perform comparably, as all the designs suffer from a fundamental limitation: To obtain a smooth dispersion curve on reflection, one must match accurately the front section of the mirror to the medium of incidence. Recently proposed implementations of CMs overcame this limitation by preventing the beams reflected at the front interface from interfering with the useful beam reflected within the multilayer.[17,18] Drawing on this concept, mirrors that provide high reflectance and accurate dispersion control over a full optical octave were demonstrated.[19,20] Mirrors of this kind manufactured by optically attaching a thin glass wedge to the top of a multilayer structure have been called tilted-front-interface chirped mirrors.[20]

Although the recently found improved implementations and the analytical predesign methods were essential for enhancing the performance of CMs, they did not obviate the need for efficient computer optimization techniques. Two major trends in CM development can currently be identified: the quest for mirror designs with even larger bandwidths, which would permit the generation of nearly single-cycle pulses for scientific applications, and the search for robust CM designs that can be reliably manufactured with a cheap coating technology and thus would be suitable for industrial production. These needs triggered the development of the optimization techniques described in this paper.

In Section 2 we describe an efficient global optimization algorithm that has been successfully tested for CM design. One of the main problems with optimization of multilayer dielectric coatings is that the merit function usually has many local extrema,[21] so there is always a risk that a local optimization routine will get stuck far away from the best extremum. Global optimization was recognized long ago as a solution to this problem. There have been many successful attempts to apply global optimization algorithms to optical multilayer structures,[21–23] but there is one common problem: All known global optimization algorithms demand a large number of evaluations of the merit function, and the optimiza-

V. Yakovlev (vladislav.iakovlev@tuwien.ac.at) is with the Institute of Automation and Electrometry, Novosibirsk, Russia, and the Institute of Photonics, Vienna University of Technology, Gusshausstrasse 27, A-1040 Vienna, Austria. G. Tempea is with the Institute of Photonics, Vienna Institute of Technology, Gusshausstrasse 27, A-1040 Vienna, Austria.

Received 18 February 2002; revised manuscript received 18 July 2002.
0003-6935/02/306514-07$15.00/0

tion process is thus often extremely time consuming. Here we show how the so-called memetic algorithm[24] can be utilized for relatively fast global optimization of chirped mirrors.

Furthermore, we present an algorithm that can improve the robustness of CM designs. The group-delay dispersion (GDD) of chirped mirrors is highly sensitive to small discrepancies between the layer thicknesses of a calculated design and those of the manufactured mirror. This effect becomes increasingly more pronounced as the required spectral range gets broader. We show how the sensitivity of CMs to manufacturing errors can be substantially reduced with the aid of the so-called stochastic quasi-gradient optimization.[25,26]

2. Memetic Algorithm

The term "memetic algorithm" is used to describe the class of evolutionary algorithms in which local search plays a significant role.[24] Here we show how the memetic algorithm (which can be considered as an improved genetic algorithm) can be utilized for relatively fast global optimization of chirped mirrors. Usually a type of genetic algorithm (a particular type of evolutionary algorithm) is used in conjunction with local optimization. Employing terminology that is specific to evolutionary computation, we call a set of trial designs a population, and by the term "member of the population" we shall mean a particular design from this set. As in any genetic algorithm the memetic optimization process goes from population to population by means of crossover and mutation operators. The difference between the genetic and the memetic algorithms is that, in the latter, each newly constructed member is locally optimized before it is included in the new population; i.e., by constructing a new set of trial designs from the old set we locally optimize each design that is going to be added to the new set. This local optimization does not necessarily have to lead to the ultimate extremum: The most important thing is that it be fast but able to considerably improve designs, bringing them closer to local extrema. The effectiveness of a memetic algorithm depends to a large extent on the effectiveness of the local optimization, which we would call partial refinement. We show now how a fast partial refinement of CMs can be implemented.

The properties of multilayer dielectric structures are calculated by the transfer matrix method[27] (see Appendix A). To calculate the reflectivity of a mirror consisting of n layers at a particular wavelength, one has to calculate the product of n transfer matrices M_k. This is the most time-consuming part of any CM optimization routine. The calculations can be made faster when the merit function has to be evaluated for several designs that differ from some particular previously characterized design by only one layer. Let us define n matrices L_k,

$$L_k = \begin{cases} 1 & k = 1 \\ \prod_{i=1}^{k-1} M_i & 1 < k \le n \end{cases}, \tag{1}$$

and n matrices R_k,

$$R_k = \begin{cases} 1 & k = n \\ \prod_{i=k+1}^{n} M_i & 1 \le k < n \end{cases}. \tag{2}$$

When all the matrices are calculated, one can quickly find the reflectivity in the case when only the kth layer is changed, because only one single matrix M_k has to be calculated and only two matrix multiplications instead of $n - 1$ have to be performed:

$$\prod_{i=1}^{n} M_i = M_1 M_2 \ldots M_{k-1} M_k M_{k+1} \ldots M_n = L_k M_k R_k. \tag{3}$$

This opens a way for quick evaluation of the gradient of the merit function, quick optimization of one particular layer, and so on. There are a few ways to implement partial refinement by using this method. For example, one can calculate the gradient and optimize the mirror along this direction (hill climbing) or optimize each layer one by one.

Many modifications of the memetic algorithm can be developed. Their comparison goes beyond the scope of this paper; we describe here a version of the algorithm that has proved to be useful for CM optimization. The merit function was defined such that lower values correspond to better designs; thus we aim to minimize the merit function. The initial population (the set of initial designs) is constructed randomly (except, maybe, a few designs that can be obtained from the analytical theory[16] or from previous optimizations). All the designs have the same number of layers made from the same materials; only the layer thicknesses are different. On each iteration two members of the population are randomly chosen according to their values of the merit function: The lower the value is, the more chances there are that the member will be selected (rank selection was used). Then a new member is formed by the crossover: A part of its layers is taken from the first design, the rest is taken from the second design, and together they form the new member; the sizes of the parts are determined randomly each time. To avoid congestion of designs in a small number of local minima we make a simple heuristic check on whether the newly formed member is in the same local minimum as one of the existing members: For each of the designs we check whether the maximal difference between layer thicknesses of the selected and the newly constructed designs is smaller than 10 nm. If it is, we try to move the designs toward each other step by step (considering each design as a point in the multidimensional parametric space), allowing only those steps that decrease the value of the merit function. If we manage to decrease the distance between them to zero, one of the designs is excluded from the population. Finally we perform a partial refinement of the new member and form a new population by adding the member to the old population. In a last step we kill the worst member of the new population if the population size exceeds the prescribed limit.

When the global optimization is completed (after the specified number of iterations is reached or the convergence criterion is fulfilled), a few best members are chosen and further optimized with the aid of a full-power local optimization (conjugate gradients algorithm,[28] for example).

3. Stochastic Quasi-Gradient Algorithm

Any manufacturing process introduces systematic and random errors into the optical thicknesses of layers. The systematic errors (e.g., a uniform increase in the thicknesses of all layers) are usually not critical for chirped mirrors, because they only slightly change the reflectance and dispersion characteristics or shift them in the spectral domain, but the random errors pose a serious problem, causing large oscillations of the mirror's GDD. The global optimization algorithm described in Section 2 leads to a set of designs of comparable quality. In analyzing these designs, one may notice that they have different sensitivities to such random perturbations of layer thicknesses. This raises the question: Is it possible to optimize this sensitivity? There have been attempts to do so; for example, sensitivity optimization is included in commercial software[29] (where it is known to work quite slowly). Still, as far as we know, nobody has yet reported how algorithms of optimization of noisy functions can be applied to the design of chirped mirrors.

Why should we make the merit function noisy and how can we do it? If one were able to evaluate the merit function for each of the trial designs by manufacturing the corresponding mirror and measuring its properties, the function would look as if some noise were added to it. To find a minimum of such a function would mean to find a point where its average value is minimal, and this solution would obviously give us the most robust design. We can simulate this noise, calculating the merit function each time not exactly for the given layer thicknesses but for randomly perturbed thicknesses for which the random perturbation simulates the perturbations introduced by the manufacturing process.

The classic optimization algorithms are barely applicable when the function values are corrupted with noise, because they rely on determination of the merit function. One of the approaches that are suitable for this case is called stochastic quasi-gradient methods.[25] The underlying idea is simple: At each iteration we estimate the gradient of the merit function and make a step in the direction opposite the gradient. It resembles the steepest-descent method, but the stochastic nature of the merit function requires specialized techniques to be able to estimate the gradient and places certain restrictions on the sequence of step sizes. If all the necessary mathematical conditions (thoroughly described in Ref. 25) are met, the algorithm is guaranteed to converge to a stationary point with probability 1, though the convergence may be slow. One of the advantages of this method is that the sensitivity does not have to be explicitly calculated to be optimized. The algorithm that we implemented is analogous to the one used for feedback control of adaptive optics,[26] for which it is termed parallel stochastic perturbative gradient descent. To estimate the gradient of the merit function we used the simultaneous perturbation approximation.[30] Let $(x_1, x_2, \ldots, x_n)$ be a vector of layer thicknesses. The noisy merit function $\tilde{J}(x_1, x_2, \ldots, x_n)$ is connected to the exact merit function $J(x_1, x_2, \ldots, x_n)$ by application of random technological perturbations to the layer thicknesses:

$$\tilde{J}(x_1, x_2, \ldots, x_n) = J[T_1(x_1), T_2(x_2), \ldots, T_n(x_n)]. \tag{4}$$

To estimate the gradient of this function we introduce m vectors of n mutually independent mean-zero random variables $(p_{k1}, p_{k2}, \ldots, p_{kn})$, $k = 1, \ldots, m$. If the variables satisfy certain mathematical conditions,[30] the most important of which is that $E|p_{ki}{}^{-1}|$ or some higher inverse moment of p_{ki} must be bounded, then the gradient of $\tilde{J}(x_1, x_2, \ldots, x_n)$ may be estimated from $2m$ measurements:

$$\nabla\tilde{J} = \frac{1}{m}\sum_{k=1}^{m}\begin{pmatrix} \dfrac{\tilde{J}_k{}^{+} - \tilde{J}_k{}^{-}}{2c_k p_{k1}} \\ \vdots \\ \dfrac{\tilde{J}_k{}^{+} - \tilde{J}_k{}^{-}}{2c_k p_{kn}} \end{pmatrix}, \tag{5}$$

where

$$\tilde{J}_k{}^{+} = J[T_1(x_1 + c_k p_{k1}), T_2(x_2 + c_k p_{k2}), \ldots, T_n(x_n + c_k p_{kn})],$$
$$\tilde{J}_k{}^{-} = J[T_1(x_1 - c_k p_{k1}), T_2(x_2 - c_k p_{k2}), \ldots, T_n(x_n - c_k p_{kn})],$$

and c_k is a positive scalar that determines the step size. It is common to take the random variables p_{ki} to be symmetrically Bernoulli distributed. A larger number of iterations m provides a more accurate estimation of the gradient, improving the convergence of the algorithm, but requires more evaluations of the merit function, increasing the time necessary for one iteration. A compromise should be found that will produce the best performance. The remarkable feature of this algorithm is that the optimal value of m is often much smaller than the number of variables n, which makes the algorithm faster than the one based on the classic estimation of the gradient.

4. Results

Global optimization outperforms local optimization algorithms, particularly if no good starting design is available. Despite progress in analytical design techniques for chirped mirrors, for several dispersive mirror design problems no direct synthesis is currently possible. One of these problems is the design of dichroic CMs (input couplers) used in femtosecond lasers, where they exhibit high reflectance and controlled negative dispersion over the fluorescence spectrum of the active laser medium and high trans-

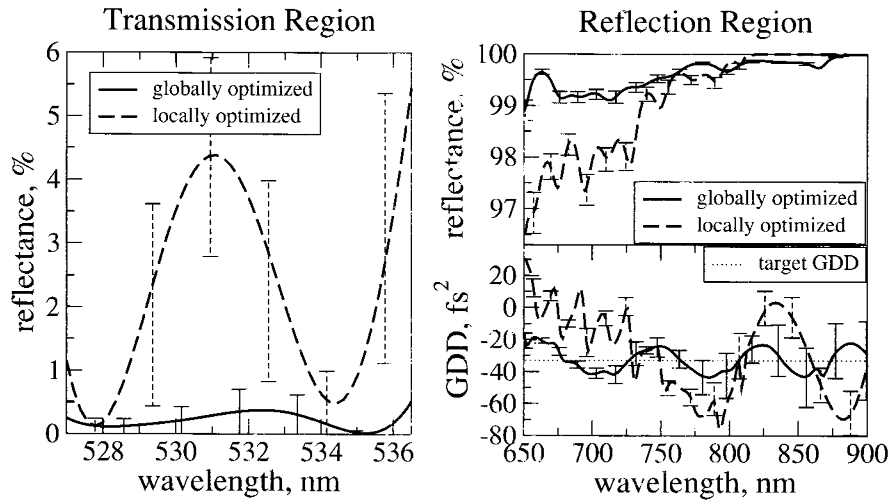

Fig. 1. Comparison of an input-coupler design found by the memetic algorithm (solid curves) with that obtained by local optimization using the conjugate-gradient algorithm (dashed curves). In both cases the starting design was obtained by the double-chirping method. Large transmittance in the wavelength range 527–537 nm as well as high reflectance in the range 650–900 nm together with a target dispersion of -35 fs^2 were required. The error bars correspond to perturbation of layer thicknesses with a standard deviation of 1 nm.

mittance over the relatively narrow wavelength range at which the medium efficiently absorbs the pump radiation. Although double-chirped mirrors inherently tend to exhibit large transmittance at wavelengths below the high-reflectance range, parameters such as the position, the width, and the losses of the high-transmittance band cannot be explicitly taken into account by analytical predesign; thus the use of a global optimization technique is called for.

To demonstrate the powerfulness of the memetic algorithm we present a CM designed to be used as an input coupler of a Ti:sapphire oscillator (Fig. 1) that has a high reflectance and a nearly constant GDD in the range 650–900 nm (the target GDD was set to -35 fs^2) and transmission as high as possible in the range 527–537 nm. The mirror consists of 60 layers made from two materials: the low-refractive-index material is SiO_2, and the refractive index of the second material is 2.25 at 800 nm. The mirror was designed, first, with the conjugate gradients algorithm starting from a double-chirped design provided by the analytical theory[16] and, second, by use of the memetic algorithm starting from the same initial stack formula. The population for the memetic algorithm consisted of 100 members; 3000 iterations of the algorithm were performed. As can be seen from Fig. 1, the global optimization found a better solution than the local refinement.

The memetic algorithm has also proved to be efficient in the design of ultrabroadband, octave-spanning chirped mirrors such as the one depicted in Fig. 2. The mirror consists of 70 layers made from the same materials as the input coupler presented above. The incidence medium for this design (fused silica) has a refractive index close to that of the low-refractive-index material; this is the so-called "back-side-coated mirror"[19] or "tilted front-interface"[20] approach—a powerful technique for suppressing the GDD oscillations and increasing the mirrors' bandwidth. To design such a broadband mirror we used a merit function based on analysis of the reflection of a trial pulse to properly weight the reflectance against GDD during optimization.

As stated above, without the partial refinement the memetic algorithm that we implemented would be just a genetic algorithm. The effectiveness of the memetic approach is demonstrated by a comparison of optimizations with and without the partial refinement. Such a comparison is shown in Fig. 3. We cannot compare the algorithms in terms of the number of merit function evaluations because the memetic algorithm uses two different ways to calculate the function, so we compare them in terms of the processor time required by the optimization (a PC with a 600-MHz Athlon processor was used). Although the partial refinement slows down the optimization process at the initial stage, it allows the memetic algorithm to outperform its genetic counterpart on the long-term run.

A result of the sensitivity optimization is shown in Fig. 4. A mirror consisting of 49 layers was optimized in the wavelength range 650–900 nm for normal incidence of light from the air. The design obtained was further optimized with the aid of a simple implementation of the stochastic quasi-gradient algorithm, in which the step size was kept constant (equal to 0.5 nm), the random variables p_{ki} were allowed to take only values $+1$ and -1 with equal probabilities, c_k were chosen to be equal to 2 nm; we performed 10,000 iterations to get the final design. The layer thicknesses before and after the sensitivity optimization are compared in Table 1. Two major assumptions were made regarding the error distribution of layer thicknesses introduced by the manufacturing process: First, the perturbations were

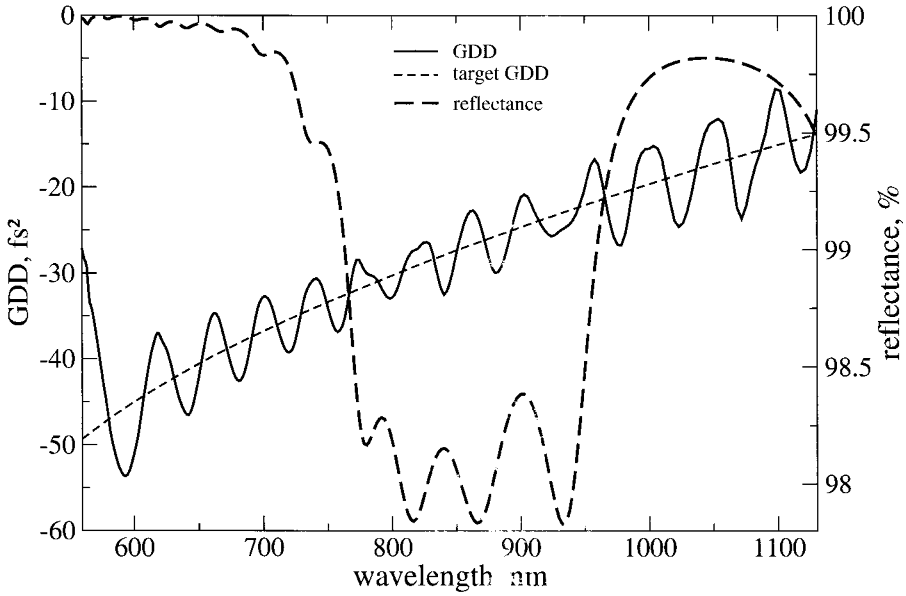

Fig. 2. Example of a broadband TFI chirped mirror designed for the wavelength range 560–1130 nm. The global and the local optimization produced designs of comparable quality in this case, probably because the initial design was close to the global optimum. The mirror consists of 70 layers, the incidence medium is fused silica, the angle of incidence is 5°, and the light is p polarized. Neither the reflectance of the fused-silica wedge to air nor the dispersion of the wedged plate was taken into account in the figure.

simulated by an additive noise that affected all the layers in the same way, and, second, the distribution was taken to be Gaussian with a standard deviation of 1 nm. These approximations agree qualitatively with the experimental measurements of available mirrors, though we cannot rigorously prove their validity. It was observed that the sensitivity optimization substantially shortens the error bars. Although it tends to decrease the quality of the ideally manufactured design (with the layer thicknesses exactly equal to the prescribed ones), this effect is overcompensated for by the improved design robustness, provided that the technological perturbations have been well enough approximated.

5. Conclusions

Using the memetic approach to global optimization, we constructed an algorithm that showed good results when it was applied to chirped-mirror design. The key component of this algorithm is the fast partial refinement of trial designs. The stochastic quasi-gradient algorithm can be utilized to reduce the mirror's sensitivity to the random perturbations of layer thicknesses that are introduced by any manufacturing process. We do not see any obstacle to

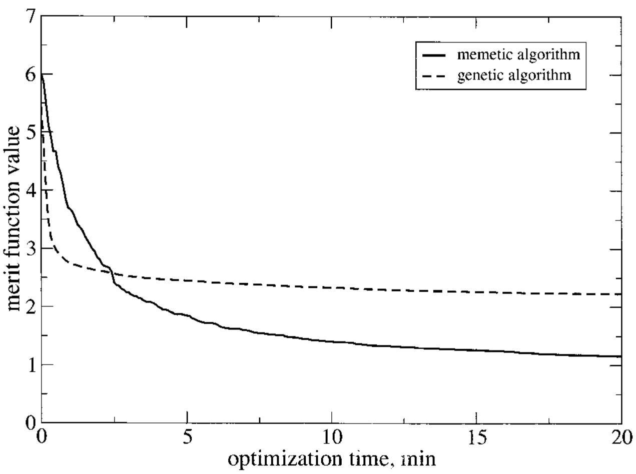

Fig. 3. Comparison of the memetic algorithm with the genetic algorithm formed by exclusion of the partial refinement. The y axis is the value of the merit function of the best member in the population. Each curve is the result of averaging over 11 runs of the optimization code.

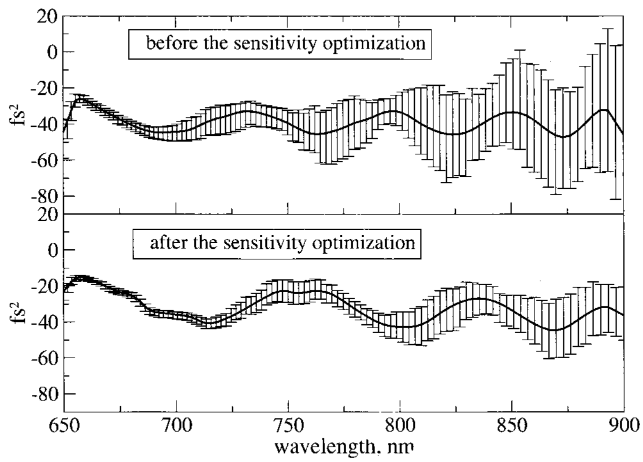

Fig. 4. Result of sensitivity optimization of a modest-bandwidth chirped mirror. The mirror consists of 49 layers and the incidence medium is air; normal incidence. The error bars were obtained for a standard deviation of technological errors equal to 1 nm.

Table 1. Comparison of Layer Thicknesses Before (d_1) and After (d_2) Sensitivity Optimization by the Stochastic Quasi-Gradient Algorithm

Layer Number	Material	d_1 (nm)	d_2 (nm)
1	SiO_2	186.8	190.7
2	TiO_2	136.9	137.6
3	SiO_2	182.2	179.6
4	TiO_2	292.4	296.4
5	SiO_2	282.2	282.4
6	TiO_2	126.6	135.2
7	SiO_2	184.5	180.1
8	TiO_2	268.0	268.4
9	SiO_2	148.3	150.5
10	TiO_2	97.9	100.0
11	SiO_2	197.5	179.0
12	TiO_2	112.4	99.7
13	SiO_2	114.2	118.3
14	TiO_2	90.8	98.3
15	SiO_2	173.6	186.7
16	TiO_2	85.8	83.7
17	SiO_2	150.0	141.5
18	TiO_2	94.8	94.6
19	SiO_2	150.7	158.4
20	TiO_2	102.1	109.6
21	SiO_2	138.5	138.9
22	TiO_2	86.6	92.3
23	SiO_2	135.1	139.5
24	TiO_2	81.0	105.4
25	SiO_2	132.2	143.1
26	TiO_2	79.7	97.1
27	SiO_2	130.7	142.2
28	TiO_2	79.1	87.7
29	SiO_2	129.9	133.4
30	TiO_2	78.9	81.5
31	SiO_2	129.4	137.3
32	TiO_2	79.2	75.4
33	SiO_2	126.2	130.0
34	TiO_2	67.9	86.1
35	SiO_2	110.8	124.0
36	Ta_2O_5	64.8	87.7
37	SiO_2	117.0	117.1
38	Ta_2O_5	88.7	79.7
39	SiO_2	131.7	124.9
40	Ta_2O_5	81.9	86.8
41	SiO_2	76.6	94.3
42	Ta_2O_5	51.6	66.1
43	SiO_2	134.1	129.6
44	Ta_2O_5	90.1	74.9
45	SiO_2	127.7	127.3
46	Ta_2O_5	11.9	32.9
47	SiO_2	130.1	117.8
48	Ta_2O_5	94.1	80.9
49	SiO_2	140.3	159.7

applying these algorithms to other kinds of optical multilayer coatings as well.

Appendix A. The Transfer Matrix Formalism

The Fresnel coefficient of reflectivity r of a thin-film multilayer structure consisting of n layers is given by

$$r = \frac{\eta_0 - Y}{\eta_0 + Y}, \tag{A1}$$

where

$$Y = C/B, \begin{pmatrix} B \\ C \end{pmatrix} = \left\{ \prod_{k=1}^{n} M_k \right\} \begin{pmatrix} 1 \\ \eta_{n+1} \end{pmatrix}, \tag{A2}$$

$$M_k = \begin{bmatrix} \cos \delta_k & (i \sin \delta_k)/\eta_k \\ i\eta_k \sin \delta_k & \cos \delta_k \end{bmatrix}, \tag{A3}$$

$$\delta_k = \frac{2\pi N_k d_k \cos \theta_k}{\lambda}, \tag{A4}$$

$$\eta_k = \begin{cases} N_k \cos \theta_k & \text{TE waves } (s \text{ waves}) \\ N_k / \cos \theta_k & \text{TM waves } (p \text{ waves}) \end{cases}, \tag{A5}$$

$$N_0 \sin \theta_0 = N_k \sin \theta_k, \tag{A6}$$

where θ_0 is the angle of incidence, d_k is the physical thickness of the kth layer, and N_k is its complex refractive index (N_0 is the index of the incidence medium and N_{n+1} is the index of the exit medium). The phase shift on reflection and the reflectance are given by $\arg(r)$ and $|r|^2$, respectively. Refractive indices N_k depend on wavelength λ, and this determines the dispersion properties of the structure.

The authors gratefully acknowledge helpful discussions with F. Krausz and A. Apolonski. This research was sponsored by the Austrian Science Fund within the framework of grant Y44-PHY as well as by the Wiener Wirtschaftsförderungsfonds.

References

1. R. Szipöcs, K. Ferencz, Ch. Spielmann, and F. Krausz, "Chirped multilayer coatings for broadband dispersion control in femtosecond lasers," Opt. Lett. **19,** 201–203 (1994).
2. A. Rundquist, C. Durfee, Z. Chang, G. Taft, E. Zeek, S. Backus, M. M. Murnane, H. C. Kapteyn, I. Christov, and V. Stoev, "Ultrafast laser and amplifier sources," Appl. Phys. B **65,** 161–174 (1997).
3. A. Stingl, M. Lenzner, Ch. Spielmann, F. Krausz, and R. Szipöcs, "Sub-10-fs mirror-dispersion-controlled Ti:sapphire laser," Opt. Lett. **20,** 602–604 (1995).
4. D. H. Sutter, I. D. Jung, F. X. Kärtner, N. Matuschek, F. Morier-Genoud, V. Schreuer, M. Tilsch, T. Tschudi, and U. Keller, "Self-starting 6.5-fs pulses from a Ti:sapphire laser using a SESAM and double-chirped mirrors," IEEE J. Sel. Top. Quantum Electron. **4,** 169–178 (1998).
5. U. Morgner, F. X. Kärtner, S. H. Cho, Y. Chen, H. A. Haus, J. G. Fujimoto, E. P. Ippen, V. Scheuer, G. Angelow, and T. Tschudi, "Sub-two-cycle pulses from a Kerr-lens mode-locked Ti:sapphire laser," Opt. Lett. **24,** 411–413 (1999).
6. R. Ell, U. Morgner, F. X. Kärtner, J. G. Fujimoto, E. P. Ippen, V. Scheuer, G. Angelow, T. Tschudi, M. J. Lederer, A. Boiko, and B. Luther-Davies, "Generation of 5-fs pulses and octave-spanning spectra directly from a Ti:sapphire laser," Opt. Lett. **26,** 373–375 (2001).
7. M. Nisoli, S. Stagira, S. de Silvestri, O. Svelto, S. Sartania, Z. Cheng, G. Tempea, Ch. Spielmann, and F. Krausz, "Toward a terawatt-scale sub-10-fs laser technology," IEEE J. Sel. Top. Quantum Electron. **4,** 414–420 (1998).
8. J. K. Ranka, R. S. Windeler, and A. J. Stentz, "Visible continuum generation in air–silica microstructure optical fibers with anomalous dispersion at 800 nm," Opt. Lett. **25,** 25–27 (2000).
9. L. Xu, W. M. Kimmel, P. O'Shea, R. Trebino, J. K. Ranka, R. S.

Windeler, and A. J. Stentz, "Measuring the intensity and phase of an ultrabroadband continuum," in *Ultrafast Phenomena XII*, T. Elsässer, S. Mukamel, M. M. Murnane, and N. F. Scherer, eds. (Springer-Verlag, Berlin, 2000), pp. 129–131.
10. Z. Cheng, G. Tempea, T. Brabec, K. Ferencz, Ch. Spielmann, and F. Krausz, "Generation of intense diffraction-limited white light and 4-fs pulses," in *Ultrafast Phenomena XI*, T. Elsässer, J. G. Fujimoto, D. A. Wiesma, and W. Zinth, eds. (Springer-Verlag, Berlin, 1998).
11. A. Baltuška, Z. Wei, M. S. Pshenichnikov, D. A. Wiersma, and R. Szipöcs, "All-solid-state cavity-dumped sub-5-fs laser," Appl. Phys. B **65,** 175–188 (1997).
12. A. Apolonski, A. Poppe, G. Tempea, Ch. Spielmann, Th. Udem, R. Holzwarth, T. W. Hänsch, and F. Krausz, "Controlling the phase evolution of few-cycle light pulses," Phys. Rev. Lett. **85,** 740–743 (2000).
13. F. X. Kärtner, N. Matuscheck, T. Tschibili, U. Keller, H. A. Haus, C. Heine, R. Morf, V. Scheuer, M. Tilsch, and T. Tschudi, "Design and fabrication of double-chirped mirrors," Opt. Lett. **22,** 831–833 (1997).
14. R. Szipöcs and A. Köhazi-Kis, "Theory and design of chirped dielectric laser mirrors," Appl. Phys. B **65,** 115–135 (1997).
15. G. Tempea, F. Krausz, Ch. Spielmann, and K. Ferencz, "Dispersion control over 150 THz with chirped dielectric mirrors," IEEE J. Sel. Top. Quantum Electron. **4,** 193–196 (1998).
16. N. Matuschek, F. X. Kärtner, and U. Keller, "Analytical design of double-chirped mirrors with custom-tailored dispersion characteristics," IEEE J. Quantum Electron. **35,** 129–137 (1999).
17. G. Tempea and F. Krausz, "Dispersion management over one octave with tilted-front-interface chirped mirrors," presented at the Twelfth International Conference on Ultrafast Phenomena, Charleston, S.C., 9–13 July 2000.
18. L. Gallmann, N. Matuschek, D. H. Sutter, V. Scheuer, G. Angelow, T. Tschudi, G. Steinmeyer, and U. Keller, "Smooth dispersion compensation: novel chirped mirrors with suppressed dispersion oscillations," in *Ultrafast Phenomena XII*, T. Elsässer, S. Mukamel, M. M. Murnane, and N. F. Scherer, eds., Springer-Verlag, Berlin, 2000), pp. 62–64.
19. N. Matuschek, L. Gallmann, D. H. Sutter, G. Steinmeyer, and U. Keller, "Back-side-coated chirped mirrors with ultra-smooth broadband dispersion characteristics," Appl. Phys. B **71,** 509–522 (2000).
20. G. Tempea, V. Yakovlev, B. Bacovic, and F. Krausz, "Tilted-front-interface chirped mirrors," J. Opt. Soc. Am. B **18,** 1747–1750 (2001).
21. T. Boudet, P. Chaton, L. Herault, G. Gonon, L. Jouanet, and P. Keller, "Thin-film designs by simulated annealing," Appl. Opt. **35,** 6219–6226 (1996).
22. J. A. Dobrowolski and R. A. Kemp, "Refinement of optical multilayer systems with different optimization procedures," Appl. Opt. **29,** 2876–2893 (1990).
23. E. Michielssen, S. Ranjithan, and R. Mittra, "Optimal multilayer filter design using real coded genetic algorithms," IEE Proc. J **139,** 413–420 (1992).
24. N. J. Radcliffe and P. D. Surry, "Formal memetic algorithms," in *Evolutionary Computing*, T. Fogarty, ed., proceedings of an AISB workshop (Springer-Verlag, Berlin, 1994), pp. 1–14.
25. Yu. Ermoliev and R. J.-B. Wets, eds. *Numerical Techniques for Stochastic Optimization* (Springer-Verlag, Berlin, 1988).
26. M. A. Vorontsov, G. W. Carhart, M. Cohen, and G. Cauwenberghs, "Adaptive optics based on analog parallel stochastic optimization: analysis and experimental demonstration," J. Opt. Soc. Am. A **17,** 1440–1453 (2000).
27. H. A. Macleod, "Basic theory," in *Thin-Film Optical Filters*, W. T. Welford and S. Laurenson, eds. (Institute of Physics Publishing, Bristol, UK, 2001), pp. 40–41.
28. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, "Conjugate gradient methods in multidimensions," in *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed., W. H. Press, ed. (Cambridge U. Press, Cambridge, 1995), pp. 420–425.
29. TFCalc, Thin-Film design software for windows, V. 3.0 (http://www.sspectra.com).
30. J. C. Spall, "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," IEEE Trans. Automatic Control **37,** 332–341 (1992).