

The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

JCAP07(2011)034

(<http://iopscience.iop.org/1475-7516/2011/07/034>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 130.226.179.254

The article was downloaded on 19/05/2012 at 20:06

Please note that [terms and conditions apply](#).

The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes

Diego Blas,^a Julien Lesgourgues^{a,b,c} and Thomas Tram^{b,d}

^aInstitut de Théorie des Phénomènes Physiques, École Polytechnique Fédérale de Lausanne, CH-1015, Lausanne, Switzerland

^bCERN, Theory Division, CH-1211 Geneva 23, Switzerland

^cLAPTh (CNRS — Université de Savoie), BP 110, F-74941 Annecy-le-Vieux Cedex, France

^dDepartment of Physics and Astronomy, University of Aarhus, DK-8000 Aarhus C, Denmark

E-mail: diego.blas@epfl.ch, julien.lesgourgues@cern.ch, tram@phys.au.dk

Received May 17, 2011

Revised June 21, 2011

Accepted July 7, 2011

Published July 22, 2011

Abstract. Boltzmann codes are used extensively by several groups for constraining cosmological parameters with Cosmic Microwave Background and Large Scale Structure data. This activity is computationally expensive, since a typical project requires from 10^4 to 10^5 Boltzmann code executions. The newly released code CLASS (Cosmic Linear Anisotropy Solving System) incorporates improved approximation schemes leading to a simultaneous gain in speed and precision. We describe here the three approximations used by CLASS for basic Λ CDM models, namely: a baryon-photon tight-coupling approximation which can be set to first order, second order or to a compromise between the two; an ultra-relativistic fluid approximation which had not been implemented in public distributions before; and finally a radiation streaming approximation taking reionisation into account.

Keywords: CMBR theory, power spectrum

ArXiv ePrint: [1104.2933](https://arxiv.org/abs/1104.2933)

Contents

1	Introduction	1
2	Tight Coupling Approximation (TCA)	2
2.1	Full equations	4
2.2	TCA equations	5
2.3	Perturbative expansion	5
2.4	Second-order approximation	6
2.5	Implementation of various schemes in CLASS	9
2.6	Comparison at the level of perturbations	9
2.7	Comparison at the level of temperature/polarisation spectra	11
3	Ultra-relativistic Fluid Approximation (UFA)	14
3.1	Truncation of the Boltzmann hierarchy	14
3.2	Sub-Hubble fluid approximation	14
3.3	Fluid equations	15
3.4	Alternative schemes	16
3.5	Comparison at the level of perturbations	16
3.6	Comparison at the level of CMB/matter power spectrum	16
4	Radiation Streaming Approximation (RSA)	17
4.1	Relativistic relics (massless neutrinos)	19
4.2	Photons	20
4.3	Summary of RSA equations	21
4.4	Comparison at the level of perturbations	21
4.5	Comparison at the level of temperature/polarisation spectra	22
5	Conclusions	24
A	Stiff integrator	25
B	Derivation of fluid equations for ultra-relativistic relics	27
B.1	Formal solution	27
B.2	Sub-Hubble approximation	29
B.3	Exact truncation formula	30
B.4	Sub-Hubble truncation formula	30

1 Introduction

Boltzmann codes have experienced considerable improvements in terms of precision and speed with respect to the pioneering COSMICS package [1]. In each new public code (CMB-FAST [2], CAMB [3], CMBEASY [4]), several sophisticated optimisation methods and approximation schemes have been introduced. Efforts on this side keeps being justified for two reasons. On the one hand, we need to fit data with higher and higher precision. For instance, the analysis of Planck data requires much more accurate theoretical predictions than

for WMAP [5]. On the other hand, a growing number of cosmologists are interested in fitting cosmological data with several extensions of the minimal cosmological model, in order to probe new physics. This requires running parameter extraction algorithms on computer clusters, often involving 10^4 or 10^5 Boltzmann code executions (for each new model or new combination of data sets). Hence, any way to speed up Boltzmann codes without losing precision is useful.

In front of such needs, a new code, the Cosmic Linear Anisotropy Solving System (CLASS) [6], has just been released.¹ The goal of this project is not just to improve speed and precision, but also to provide a flexible and user-friendly code that can be easily generalized to non-minimal cosmological models. In this paper, we do not discuss flexibility issues and only concentrate on the improved approximation schemes used by CLASS, in the strict context of minimal Λ CDM cosmology. Extensions requiring extra approximations may be introduced and discussed case by case in the future. In a companion paper [7], we already discuss the approximation specific to massive neutrinos and non-cold dark matter relics. A comparison between the power spectra obtained by CAMB and CLASS for the minimal Λ CDM model, as well as estimates of the relative speed of the two codes, is presented in [8].

The next three sections describe: a baryon-photon tight-coupling approximation which can be set to first order, second order or to a compromise between the two (section 2); an ultra-relativistic fluid approximation which had not been implemented in public distributions of Boltzmann codes before (section 3); and a radiation streaming approximation consistently including reionisation (section 4). Appendix A describes the stiff integrator which can be used by CLASS as an alternative to the Runge-Kutta integrator: without this integrator, it would have been essentially impossible to launch test runs with no approximation schemes, to evaluate the error induced by these schemes; moreover, this integrator gives better performances even in the presence of approximations, and is set to be the default integrator in CLASS. We summarize our full approximation landscape in section 5.

Note that throughout this paper, when discussing CAMB and CLASS, we refer to the versions available at the time of preparing this manuscript, i.e. the January 2011 version of CAMB and v1.1 of CLASS.

2 Tight Coupling Approximation (TCA)

Before recombination, when the opacity $an_e\sigma_T$ is very large, the equations governing the tightly coupled baryon-photon fluid form a stiff system. Indeed, the opacity defines a conformal time scale of interaction $\tau_c \equiv (an_e\sigma_T)^{-1}$ considerably smaller than that on which most of the modes actually evolve, namely $\tau_H = a/a'$ for super-Hubble scales and $\tau_k = 1/k$ for sub-Hubble scales (see figure 1). Standard integrators like Runge-Kutta algorithms would be very inefficient in solving such a system. This motivated Peebles & Yu to introduce a simplified system of differential equations valid in the regime of small τ_c/τ_H and τ_c/τ_k (tight coupling approximation or TCA) [9]. The overall idea is that quantities which are vanishingly small in the limit $\tau_c \rightarrow 0$ are solved perturbatively in τ_c , and these analytical expressions are used in the numerical code solving the remaining differential equations of the system.

The CLASS user can choose between two integrators for the system of linear perturbations. One of them (`ndf15`, see appendix A) is optimised for stiff equations, and shows good performances even in the tight coupling regime. However, by reducing drastically the number

¹Available at <http://class-code.net>.




Figure 1. (Left) Evolution characteristic conformal time scales in units of Mpc. Before recombination, the baryon-photon interaction time scale $\tau_c = (an_e\sigma_T)^{-1}$ is much smaller than the Hubble time scale $\tau_H = a/a'$. For each mode and in a given range, it is also smaller than the acoustic oscillation time scale $1/k$. (Right) Evolution of the product $a c_a^2$ (scale factor times baryon adiabatic sound speed), with an arbitrary normalization of the scale factor, before and during hydrogen recombination (in this model, the visibility function peaks at $\tau_{\text{rec}} = 278$ Mpc). For $\tau < 200$ Mpc, helium recombination leads to a variation of this product by approximately 10%.

of equations to integrate, any TCA scheme leads to a speed up even in presence of such an integrator.

Tight coupling equations have already been derived and improved by many authors after Peebles & Yu’s seminal paper. For instance, the TCA formulas presented in Ma & Bertschinger [10] are derived to first order in τ_c (omitting some polarisation terms which contribute to the photon shear at this order). Lewis et al. implemented the full first-order solution in CAMB [3], also relaxing Ma & Bertschinger’s assumption that $\tau_c \propto a^2$. Doran implemented some improved formulas in CMBEASY which are valid in the Newtonian gauge, and include a few contributions beyond order one [11].

As we will see below, the first-order TCA formulas provide poor approximations to the baryon-photon differential energy flux and to the photon shear at large times. This is not much of a problem if one switches to the exact² equations early enough. However, a better scheme would allow to switch off the TCA later, and to save a lot of integration time without losing precision.

Here, we will derive the full second-order TCA formulas in the synchronous gauge. While this work was in preparation, Cyr-Racine and Sigurdson published a paper on exactly the same topic [12]. We will show that the numerical results from our approach coincide with those of [12]. Another recent paper discussing the TCA beyond first order and its implementation in second order cosmological perturbation theory is [14]: this work actually proposes a systematic way to compute high-order corrections in a given model and at any order.³

²It should be intended that throughout this section, the word “exact” is intended in the sense “without using the TCA”. Obviously, our equations are never exact since they rely on a number of common and well-justified approximations, e.g.: linear perturbations, pressureless CDM, approximate expression for photon-baryon coupling, baryon pressure neglected in several equations, etc.

³As pointed out by the author of [14], this method can be implemented numerically provided that the code computing the evolution of thermodynamical variables outputs fully continuous and derivable functions of time, which is precisely the case with the CLASS version of RECFAST.

2.1 Full equations

In the following we will adopt the notation of Ma & Bertschinger [10]. The baryon perturbations will be characterized by the energy density contrast δ_b and the divergence of the fluid energy flux θ_b . These quantities satisfy the equations⁴

$$\delta'_b = -\theta_b - \frac{1}{2}h', \quad (2.1a)$$

$$\theta'_b = -\mathcal{H}\theta_b + c_s^2 k^2 \delta_b + \frac{R}{\tau_c} \Theta_{\gamma b}, \quad (2.1b)$$

with $\mathcal{H} = aH = a'/a$, where a is the scale factor, and we have defined $\Theta_{\gamma b} \equiv \theta_\gamma - \theta_b$ and $R \equiv \frac{4\bar{p}_\gamma}{3\bar{p}_b}$. The quantity $\Theta_{\gamma b}$ represents the divergence of the energy flux of the photons (θ_γ) in the frame comoving with the baryons. Its time-derivative $\Theta'_{\gamma b}$ is often referred to as the ‘‘baryon-photon slip’’. The fields h and η represent the metric perturbations. As explained in [10], the baryon pressure perturbation can be safely neglected in the continuity equation, but its Laplacian should be kept in the Euler equation (giving raise to the term $c_s^2 k^2 \delta_b$) since it affects the evolution of very small wavelengths, smaller than the baryon Jeans length. In [10] and many other references, the baryon sound speed is identified to the adiabatic sound speed

$$c_a^2 = \frac{k_B T_b}{\mu} \left(1 - \frac{1}{3} \frac{d \ln T_b}{d \ln a} \right), \quad (2.2)$$

where the evolution of the proton temperature is given by

$$T'_b = -2\mathcal{H}T_b + \frac{2\mu R}{m_e \tau_c} (T_\gamma - T_b), \quad (2.3)$$

and μ is the mean molecular weight. This approximation has been proved to be inaccurate in refs. [15, 16], but the difference is only important for computing the matter power spectrum for $k \gg 10 h \text{ Mpc}^{-1}$. The current version of CLASS (v1.1) and CAMB (from January 2011) still relies on the $c_s^2 = c_a^2$ approximation, while future versions of both codes are likely to switch to the actual sound speed calculation, as in the CAMB_source⁵ code. This issue is irrelevant for the results of this paper, which do not involve very large wavelengths. All numerical results below have been obtained using c_a^2 instead of c_s^2 , but our formulas can adequately describe the large k range, provided that a correct sound speed calculation is performed.

Like the adiabatic sound speed, the actual baryon sound speed is expected to decrease approximately as a^{-1} , except during helium and hydrogen recombination (see figure 1).

The characterization of the photon distribution requires the determination of its different multipoles δ_γ , θ_γ , σ_γ and $F_{\gamma l}$ for $l \geq 3$. These satisfy the recursive Boltzmann equations (eqs. (63) in [10]):

$$\delta'_\gamma = -\frac{4}{3}\theta_\gamma - \frac{2}{3}h', \quad (2.4a)$$

$$\theta'_\gamma = k^2 \left(\frac{1}{4}\delta_\gamma - \sigma_\gamma \right) - \frac{\Theta_{\gamma b}}{\tau_c}, \quad (2.4b)$$

$$2\sigma'_\gamma = \frac{8}{15}\theta_\gamma - \frac{3}{5}kF_{\gamma 3} + \frac{4}{15}(h' + 6\eta') - \frac{9}{5\tau_c}\sigma_\gamma + \frac{1}{10\tau_c}(G_{\gamma 0} + G_{\gamma 2}), \quad (2.4c)$$

⁴We will work in Fourier space and stick to the synchronous gauge in conformal time, which we will denote by τ . We use the prime to denote derivative with respect to conformal time.

⁵<http://camb.info/sources/>.

$$F'_{\gamma l} = \frac{k}{2l+1} [lF_{\gamma(l-1)} - (l+1)F_{\gamma(l+1)}] - \frac{1}{\tau_c} F_{\gamma l}, \quad l \geq 3 \quad (2.4d)$$

$$G'_{\gamma l} = \frac{k}{2l+1} [lG_{\gamma(l-1)} - (l+1)G_{\gamma(l+1)}] + \frac{1}{\tau_c} \left[-G_{\gamma l} + \frac{1}{2}(F_{\gamma 2} + G_{\gamma 0} + G_{\gamma 2}) \left(\delta_{l0} + \frac{\delta_{l2}}{5} \right) \right], \quad (2.4e)$$

where we remind the reader that $F_{\gamma 2} = 2\sigma_\gamma$. Finally, the previous hierarchical equations are truncated at some $l = l_{\max}$ following equation (65) in [10],

$$F'_{\gamma l_{\max}} = kF_{\gamma(l_{\max}-1)} - \frac{l_{\max}+1}{\tau} F_{\gamma l_{\max}} - \tau_c^{-1} F_{\gamma l_{\max}}. \quad (2.5)$$

2.2 TCA equations

From eqs. (2.1) and (2.4) one sees that the different time scales in the problem are τ_c , k^{-1} and the time scale of cosmological evolution \mathcal{H} . As we will show explicitly in the next section, it is possible to find a solution $(\Theta'_{\gamma b}{}^{\text{tca}}, \sigma_\gamma{}^{\text{tca}})$ for the baryon-photon slip and the photon shear in terms of δ_b , δ_γ , θ_b and θ_γ , which is valid to any desired order in the small parameter τ_c . The knowledge of this solution helps to reduce the full system of equations (2.1) and (2.4) to just four of them for the low multipoles of the distributions. More concretely, one may use eq. (2.1a) for δ_b' , and eq. (2.4a) to determine δ_γ' ; the energy fluxes are characterized by the linear combination of eqs. (2.1b) and (2.4b) in which the coupling term vanishes:

$$\theta_b' + R\theta_\gamma' = -\mathcal{H}\theta_b + c_s^2 k^2 \delta_b + Rk^2 \left(\frac{1}{4} \delta_\gamma - \sigma_\gamma{}^{\text{tca}} \right), \quad (2.6)$$

and finally $\theta_\gamma' - \theta_b' = \Theta'_{\gamma b}{}^{\text{tca}}$. As desired, this scheme allows us to get rid of any coefficient⁶ in τ_c^{-1} . For practical reasons it is customary to combine linearly the last two equations in order to eliminate θ_γ' and get an expression for θ_b' only; then θ_γ' can be found from equation (2.6). In summary, once $\Theta'_{\gamma b}{}^{\text{tca}}$ and $\sigma_\gamma{}^{\text{tca}}$ are known the goal is to solve the closed system formed by the four equations (2.1a), (2.4a), and

$$\theta_b' = -\frac{1}{(1+R)} \left(\mathcal{H}\theta_b - c_s^2 k^2 \delta_b - k^2 R \left(\frac{1}{4} \delta_\gamma - \sigma_\gamma{}^{\text{tca}} \right) + R\Theta'_{\gamma b}{}^{\text{tca}} \right), \quad (2.7a)$$

$$\theta_\gamma' = -R^{-1}(\theta_b' + \mathcal{H}\theta_b - c_s^2 k^2 \delta_b) + k^2 \left(\frac{1}{4} \delta_\gamma - \sigma_\gamma{}^{\text{tca}} \right). \quad (2.7b)$$

2.3 Perturbative expansion

The aim of this section is to find expressions for $\Theta'_{\gamma b}$ and σ_γ valid at the n^{th} order in τ_c (the zero order is trivial: both species behave as a single perfect fluid, so that $\Theta_{\gamma b}$ and all multipoles of the photons beyond δ_γ and θ_γ vanish; the first order can be found in [10]). We first multiply equations (2.4b) and (2.1b) by τ_c :

$$\tau_c \left[\theta_\gamma' - k^2 \left(\frac{1}{4} \delta_\gamma - \sigma_\gamma \right) \right] + \Theta_{\gamma b} = 0, \quad (2.8a)$$

$$\tau_c \left[-\theta_b' - \mathcal{H}\theta_b + c_s^2 k^2 \delta_b \right] + R\Theta_{\gamma b} = 0. \quad (2.8b)$$

⁶The scale τ_c appears now as a small perturbation.

To obtain a differential equation for $\Theta_{\gamma b}$ we can combine the above two equations into

$$\tau_c \left[\Theta_{\gamma b}' - \mathcal{H}\theta_b + k^2 \left(c_s^2 \delta_b - \frac{1}{4} \delta_\gamma + \sigma_\gamma \right) \right] + (1+R)\Theta_{\gamma b} = 0. \quad (2.9)$$

This equation involves the photon shear, given by the following equation (see (2.4c)):

$$\sigma_\gamma = \frac{\tau_c}{9} \left[\frac{8}{3} \theta_\gamma + \frac{4}{3} h' + 8\eta' - 10\sigma_\gamma' - 3kF_{\gamma 3} \right] + \frac{1}{18} (G_{\gamma 0} + G_{\gamma 2}). \quad (2.10)$$

Until now, all these equations are exact. In the limit of interest, both of them can be schematically written as⁷

$$\epsilon y(t)' + y(t)/f(t) + \epsilon g(t) = 0, \quad (2.11)$$

where ϵ is a small parameter. In our case, ϵ can be chosen to be $\bar{\tau}_c$, the opacity at an arbitrary time around which the expansion is performed.⁸ The perturbative solution is given by

$$y(t) = \sum_{n=1} \epsilon^n y_n(t), \quad y_1 = -fg, \quad y_{n+1} = -fy_n'. \quad (2.12)$$

Notice that for functions $f(t)$ and $g(t)$ with smooth time variations on the scale $\bar{\tau}_c$, the previous is a perfectly well defined solution. Finally, the most general solution is found by adding to the previous particular solutions the solution of the homogeneous equation

$$\epsilon y(t)' + y(t)/f(t) = 0, \quad y = Ce^{-1/\epsilon \int f^{-1} dt}. \quad (2.13)$$

Note that in our case f is always positive, which is enough to make this part of the solution suppressed very fast. Hence, the relevant solution is given by eq. (2.12), which, after absorbing the small parameter in the function $\tilde{f} \equiv \epsilon f$, reads:

$$y(t) = \sum_{n=1} \tilde{y}_n(t), \quad \tilde{y}_1 = -\tilde{f}g, \quad \tilde{y}_{n+1} = -\tilde{f}\tilde{y}_n'. \quad (2.14)$$

In terms of these functions, the powers of \tilde{f} and its derivatives represent the different orders of the approximation.

2.4 Second-order approximation

Using the previous expansion in eq. (2.9), the baryon-photon relative velocity reads at order two:

$$\Theta_{\gamma b} = \tilde{f}_\Theta (-g_\Theta + \tilde{f}'_\Theta g_\Theta + \tilde{f}_\Theta g'_\Theta) + O(\bar{\tau}_c^3), \quad (2.15)$$

with

$$\tilde{f}_\Theta = \frac{\tau_c}{1+R}, \quad g_\Theta = -\mathcal{H}\theta_b + k^2 \left(c_s^2 \delta_b - \frac{1}{4} \delta_\gamma + \sigma_\gamma \right). \quad (2.16)$$

We still need to differentiate this equation in order to get a similar approximation for the slip:

$$\Theta'_{\gamma b} = \left(\frac{\tilde{f}'_\Theta}{\tilde{f}_\Theta} \right) \Theta_{\gamma b} + \tilde{f}_\Theta (-g'_\Theta + \tilde{f}''_\Theta g_\Theta + 2\tilde{f}'_\Theta g'_\Theta + \tilde{f}_\Theta g''_\Theta) + O(\bar{\tau}_c^3). \quad (2.17)$$

⁷This is immediate for (2.9). For (2.10) it follows from (2.4e), and we will explicitly verify it shortly.

⁸In fact, the small dimensionless parameters will be $\bar{\tau}_c k$ and $\bar{\tau}_c \bar{\mathcal{H}}$, with $\bar{\mathcal{H}}$ evaluated around the same time as $\bar{\tau}_c$.

There are several ways to organise and simplify the final result. In order to write the first-order term in the same form as in the rest of the literature, we need to express g'_Θ in a very peculiar way:

$$\begin{aligned}
 g'_\Theta &= -\mathcal{H}\theta'_b - \mathcal{H}'\theta_b + k^2 \left(c_s'^2 \delta_b + c_s^2 \delta'_b - \frac{1}{4} \delta'_\gamma + \sigma'_\gamma \right) \\
 &= -2\mathcal{H}\theta'_b - (\mathcal{H}' + \mathcal{H}^2)\theta_b + k^2 \left((\mathcal{H}c_s^2 + c_s'^2)\delta_b + c_s^2 \delta'_b - \frac{1}{4} \delta'_\gamma + \sigma'_\gamma \right) + \frac{R\mathcal{H}}{\tau_c} \Theta_{\gamma b} \quad (2.18) \\
 &= 2\mathcal{H}\Theta_{\gamma b}' - \frac{a''}{a} \theta_b + k^2 \left(-\frac{\mathcal{H}}{2} \delta_\gamma + \bar{c}_s^2 \delta_b + c_s^2 \delta'_b - \frac{1}{4} \delta'_\gamma + 2\mathcal{H}\sigma_\gamma + \sigma'_\gamma \right) + \frac{(2+R)\mathcal{H}}{\tau_c} \Theta_{\gamma b}.
 \end{aligned}$$

Here we defined $\bar{c}_s^2 \equiv (\mathcal{H}c_s^2 + c_s'^2)$: this quantity would vanish if the approximation $c_s^2 \propto a^{-1}$ were valid at all times. In the second line, we used eq. (2.1b), while in the third line we used (2.4b): so these expressions for g'_Θ are all exact.

The first-order approximation for $\Theta'_{\gamma b}$ is obtained by replacing the first occurrence of g'_Θ in eq. (2.17) with the last expression of (2.18), in which we neglect the terms $2\mathcal{H}\Theta'_{\gamma b}$ and $(2\mathcal{H}\sigma_\gamma + \sigma'_\gamma)$ which represent contributions of higher order. The final result is:

$$\Theta'_{\gamma b} = \left(\frac{\tau'_c}{\tau_c} - \frac{2\mathcal{H}}{1+R} \right) \Theta_{\gamma b} - \tilde{f}_\Theta \left[-\frac{a''}{a} \theta_b + k^2 \left(-\frac{\mathcal{H}}{2} \delta_\gamma + \bar{c}_s^2 \delta_b + c_s^2 \delta'_b - \frac{1}{4} \delta'_\gamma \right) \right] + O(\bar{\tau}_c^2). \quad (2.19)$$

Note that in the previous expression we used the exact relation $R' = -\mathcal{H}R$.

For the second-order expression for the slip, we go back to eq. (2.17). We replace the first occurrence of g'_Θ by the full expression (2.18), assuming that $\Theta'_{\gamma b}$ and $(2\mathcal{H}\sigma_\gamma + \sigma'_\gamma)$ have been replaced by their first-order approximation. Finally, we replace g_Θ , g'_Θ and g''_Θ in the last three terms by their zeroth-order approximation. The final result can be written in a compact form:

$$\begin{aligned}
 \Theta'_{\gamma b} &= (1 - 2\mathcal{H}\tilde{f}_\Theta) \left\{ \left(\frac{\tau'_c}{\tau_c} - \frac{2\mathcal{H}}{1+R} \right) \Theta_{\gamma b} - \tilde{f}_\Theta \left[-\frac{a''}{a} \theta_b + k^2 \left(-\frac{\mathcal{H}}{2} \delta_\gamma + \bar{c}_s^2 \delta_b + c_s^2 \delta'_b - \frac{1}{4} \delta'_\gamma \right) \right] \right\} \\
 &\quad - \tilde{f}_\Theta k^2 (2\mathcal{H}\sigma_\gamma + \sigma'_\gamma) + \tilde{f}_\Theta [\tilde{f}_\Theta'' g_\Theta + 2\tilde{f}_\Theta' g'_\Theta + \tilde{f}_\Theta g''_\Theta] + O(\bar{\tau}_c^3). \quad (2.20)
 \end{aligned}$$

This formula requires an expression for the shear valid at order one. However, to solve equations (2.1b), (2.4b) consistently to second order, we need the expression for the shear at the corresponding order. This can be achieved using (2.10). To solve this equation let us first note that the polarisation multipoles $l = 0, 2$ obey (cf. (2.4e))

$$\begin{aligned}
 G'_{\gamma 0} &= -kG_{\gamma 1} + \tau_c^{-1} \left[-G_{\gamma 0} + \sigma_\gamma + \frac{1}{2}(G_{\gamma 0} + G_{\gamma 2}) \right], \\
 G'_{\gamma 2} &= \frac{k}{5}(2G_{\gamma 1} - 3G_{\gamma 3}) + \tau_c^{-1} \left[-G_{\gamma 2} + \frac{1}{10}(2\sigma_\gamma + G_{\gamma 0} + G_{\gamma 2}) \right], \quad (2.21)
 \end{aligned}$$

from which we see that, at first order in $\bar{\tau}_c$, $G_{\gamma 2} \sim G_{\gamma 0} \sim \sigma_\gamma$. Thus, it is consistent to consider these multipoles as $O(\bar{\tau}_c)$, and write the second order solution to (2.10) following (2.14) as

$$\begin{aligned}
 \sigma_\gamma &= \frac{\tau_c}{9} \left[\frac{8}{3} \theta_\gamma + \frac{4}{3} h' + 8\eta' \right] + \frac{1}{18} (G_{\gamma 0} + G_{\gamma 2}) \\
 &\quad - \frac{10\tau_c}{9} \frac{d}{d\tau} \left(\frac{\tau_c}{9} \left[\frac{8}{3} \theta_\gamma + \frac{4}{3} h' + 8\eta' \right] + \frac{1}{18} [G_{\gamma 0} + G_{\gamma 2}] \right) + O(\bar{\tau}_c^3), \quad (2.22)
 \end{aligned}$$

where we used the fact that $F_{\gamma 3} = O(\bar{\tau}_c^2)$. Indeed, the high photon multipoles obey at leading order (cf. (2.4d))

$$F_{\gamma l} = \frac{l\tau_c k}{2l+1} F_{\gamma(l-1)}.$$

The final step involves an evaluation of (2.21). Again, from (2.4e), one finds that

$$G_{\gamma 1} \sim G_{\gamma 3} \sim O(\bar{\tau}_c^2).$$

This allows us to find the perturbative solution to (2.21)

$$G_{\gamma 0} = \frac{5\sigma_\gamma}{2} - \frac{25}{4}\tau_c\sigma'_\gamma + O(\bar{\tau}_c^3), \quad G_{\gamma 2} = \frac{\sigma_\gamma}{2} - \frac{5}{4}\tau_c\sigma'_\gamma + O(\bar{\tau}_c^3). \quad (2.23)$$

From the previous expression and (2.22) we find

$$\sigma_\gamma = \frac{8\tau_c}{45}(2\theta_\gamma + h' + 6\eta') + O(\bar{\tau}_c^2), \quad (2.24)$$

which implies

$$\sigma'_\gamma = \frac{8\tau_c}{45}(2\theta'_\gamma + h'' + 6\eta'') + \frac{8\tau'_c}{45}(2\theta_\gamma + h' + 6\eta') + O(\bar{\tau}_c^2). \quad (2.25)$$

Finally, the shear at second order is found from (2.22) to be

$$\sigma_\gamma = \frac{8\tau_c}{45} \left[(2\theta_\gamma + h' + 6\eta') \left(1 - \frac{11\tau'_c}{6} \right) - \frac{11\tau_c}{6} (2\theta'_\gamma + h'' + 6\eta'') \right] + O(\bar{\tau}_c^3). \quad (2.26)$$

The last missing items are the zero-order expressions for g_Θ , g'_Θ and g''_Θ appearing in equation (2.20). Noticing that equation (2.7a) implies

$$\begin{aligned} \theta'_b &= \frac{1}{1+R} \left(-\mathcal{H}\theta_b + k^2 c_s^2 \delta_b + k^2 R \frac{1}{4} \delta_\gamma \right) + O(\bar{\tau}_c), \\ \theta''_b &= \frac{1}{1+R} \left((R-1)\mathcal{H}\theta'_b - \mathcal{H}'\theta_b + k^2 \left((c_s^2)' \delta_b + c_s^2 \delta'_b - \frac{R\mathcal{H}}{4} \delta_\gamma + \frac{R}{4} \delta'_\gamma \right) \right) + O(\bar{\tau}_c), \end{aligned}$$

we can write these last terms as

$$g_\Theta = -\mathcal{H}\theta_b + k^2 \left(c_s^2 \delta_b - \frac{1}{4} \delta_\gamma \right) + O(\bar{\tau}_c), \quad (2.27a)$$

$$g'_\Theta = -\mathcal{H}\theta'_b - \mathcal{H}'\theta_b + k^2 \left[(c_s^2)' \delta_b + \left(\frac{1}{3} - c_s^2 \right) \left(\theta_b + \frac{1}{2} h' \right) \right] + O(\bar{\tau}_c), \quad (2.27b)$$

$$\begin{aligned} g''_\Theta &= -\mathcal{H}\theta''_b - 2\mathcal{H}'\theta'_b - \mathcal{H}''\theta_b \\ &\quad + k^2 \left[(c_s^2)'' \delta_b - 2(c_s^2)' \left(\theta_b + \frac{1}{2} h' \right) + \left(\frac{1}{3} - c_s^2 \right) \left(\theta'_b + \frac{1}{2} h'' \right) \right] + O(\bar{\tau}_c). \end{aligned} \quad (2.27c)$$

The derivation given in [12] follows different steps, but since it is still a second-order TCA, the results should be identical under the approximation $c_s^2 \propto a^{-1}$ used in [12], at least up to terms of order three or higher. For the shear, our expressions are indeed exactly identical. For the slip, there are so many ways to write the result and so many terms involved that the comparison is not trivial. However, by coding the two formulas in CLASS and comparing the evolution of $\Theta_{\gamma b}$ in the two cases, we found that the two expressions agree very well, since numerically the difference appears to be at most of order $O(\bar{\tau}_c^3)$.

2.5 Implementation of various schemes in CLASS

We implemented various TCA schemes in CLASS, which can be chosen by setting the flag `tight_coupling_approximation` to different values. We always start integrating the wavenumbers very deep inside the tightly-coupled regime. Hence, in contrast to [12], we do not need to include terms in τ_c in the initial conditions. We always use the set of equations (2.1a), (2.4a), (2.7a), (2.7b) with different expressions for the slip $\Theta_{\gamma b}^{\text{tca}}$ and the shear $\sigma_{\gamma}^{\text{tca}}$:

1. First-order expressions (2.19), (2.24), with the approximations $\tau_c \propto a^2$ and $c_s^2 \propto a^{-1}$, used e.g. in ref. [10]. This corresponds to the scheme used by CLASS when the label `tight_coupling_approximation` is set to `first_order_MB`.
2. First-order expressions (2.19), (2.24) with the only approximation $c_s^2 \propto a^{-1}$, like in CAMB. This scheme is used by CLASS when the same flag is set to `first_order_CAMB`.
3. Exact first-order expressions (2.19), (2.24) when the flag is set to `first_order_CLASS`.
4. Second-order expressions from [12] when the flag is set to `second_order_CRS`.
5. Second-order expressions from eqs. (2.20), (2.26) for the flag `second_order_CLASS`.
6. Finally, second-order expression for the shear, but a reduced expression for the slip, involving only the leading order-two terms:

$$\Theta_{\gamma b}^{\text{tca}} = (1 - 2\mathcal{H}\tilde{f}_{\Theta}) \left\{ \left(\frac{\tau'_c}{\tau_c} - \frac{2\mathcal{H}}{1+R} \right) \Theta_{\gamma b} - \tilde{f}_{\Theta} \left[-\frac{a''}{a}\theta_b + k^2 \left(-\frac{\mathcal{H}}{2}\delta_{\gamma} + c_s^2\delta'_b - \frac{1}{4}\delta'_{\gamma} \right) \right] \right\} - \tilde{f}_{\Theta} k^2 \left[2\mathcal{H}\sigma_{\gamma} + \sigma'_{\gamma} - \left(\frac{1}{3} - c_s^2 \right) \left(\tilde{f}_{\Theta}\theta'_b + 2\tilde{f}'_{\Theta}\theta_b \right) \right]. \quad (2.28)$$

This option is taken when the flag is set to `compromise_CLASS`, and is chosen to be the default option in CLASS. Notice that the last scheme does not have a term h'' , which is advantageous from the computational point of view.⁹

To justify the compromise scheme, notice that it encapsulates the leading order in (2.27) for subhorizon modes (note that θ_b has an extra momentum dependence with respect to the other perturbations, and each time derivative adds one more power of k in this regime). Thus, from eq. (2.20) we learn that eq. (2.28) implements the leading second order correction for the modes with a big comoving momentum. It is precisely for these modes that the first order approximation fails first, which explains the success of the compromise scheme. To check that the approximation is indeed correct, we implemented this scheme for several modes k assuming Λ CDM. As illustrated in the next subsection, this scheme is nearly as good as the full second-order one, being at the same time much more compact and requiring many less floating point operations.

2.6 Comparison at the level of perturbations

In figure 2, we compare these different approximations for a fixed wave number (chosen to be $k = 10^{-2} \text{Mpc}^{-1}$), and in figure 3 for a fixed conformal time τ (chosen to be the time when the TCA is switched off in the previous example).

⁹In order to compute h'' one should use one more Einstein equation than in the standard case, and compute the δT_i^i component of the stress-energy tensor, i.e. the pressure perturbation for all species.




Figure 2. Evolution of $\Theta_{\gamma b}$ (left) and σ_{γ} (right) for the mode $k = 10^{-2} \text{Mpc}^{-1}$, using the various TCA schemes listed in section 2.5. In each case, the quantities are represented as points when the TCA is switched on, and as continuous lines of the same color when exact equations take over. The TCA is switched off at $\tau = 1 \text{Mpc}$ in the reference case, and at $\tau = 194 \text{Mpc}$ in all other cases. We show a single set of points for cases which are indistinguishable by eye, namely: `first_order_CAMB` and `first_order_CLASS`; and also, `second_order_CRS` and `second_order_CLASS`. The default scheme `compromise_CLASS` is also hardly distinguishable from the `second_order_CLASS`.




Figure 3. Comparison of $\Theta_{\gamma b}(k)$ (left) and $\sigma_{\gamma}(k)$ (right) at the time when the TCA is switched off ($\tau = 194 \text{Mpc}$ in this case), for the different schemes listed in section 2.5. All cases are compared to the full second-order scheme `second_order_CLASS`. For $\Theta_{\gamma b}$, we show a single curve for the cases `first_order_CAMB` and `first_order_CLASS`, since they are indistinguishable by eye. For σ_{γ} , all schemes using the first-order shear expression are indistinguishable; this is also true for all cases using the second-order shear expression (despite the fact that θ_{γ} and metric perturbations are slightly different in each individual case).

Scrutinising first the various first-order schemes, we see a significant difference at late time between the first two (MB and CAMB), showing that $\tau_c \propto a^{-1}$ is a bad approximation. However, there is no sizable difference between the second one (CAMB) and the third one (CLASS) in which the approximation $c_s^2 \propto a^{-1}$ is relaxed. We reach the same conclusion when comparing second-order schemes with or without the same approximation. This is not a surprise, since we are only considering scales larger than the baryonic Jeans length at any time. When studying very small wavelengths, the CLASS user is free to choose one of the TCA schemes where the full evolution of c_s^2 is automatically taken into account (namely, `first_order_CLASS`

or `second_order_CLASS`), but as mentioned before, in this limit, an accurate sound speed computation should also be implemented in order to relax the $c_s^2 = c_a^2$ approximation.

Figures 2 and 3 show that all first-order schemes provide poor approximations for the slip and the shear near the end of the tightly coupled regime, and hence, inaccurate initial conditions at the time at which the full equations are turned on. As expected, second-order schemes work much better. We find a very good agreement between `second_order_CRS` and `second_order_CLASS`: this validates both the results of [12] and our results. The residual difference is likely due to the fact that our two independent derivations lead to expressions in which some higher-order terms (of order $O(\bar{\tau}_c^3)$) appear in different ways.

These two schemes also agree to a very good extent with `compromise_CLASS`, which is much more straightforward to code, and computes the baryon-photon slip with approximately ten times less operations. In particular, with this scheme, the code does not even need to compute the quantities h'' , τ_c'' and \mathcal{H}'' , which are not so obvious to obtain with few operations and good accuracy. Hence, this method is set to be the default TCA in `CLASS`.

2.7 Comparison at the level of temperature/polarisation spectra

First, let us specify which precision parameters in `CLASS` govern the evolution of perturbations in the early universe and the TCA switching time:

- Two parameters define the time at which initial conditions are imposed during the tightly-coupled stage. Each wave-number starts being integrated (with one of the TCA schemes) as soon as one of the two conditions

$$(\tau_c/\tau_H) \geq \text{start_small_k_at_tau_c_over_tau_h}$$

or

$$(\tau_H/\tau_k) \geq \text{start_large_k_at_tau_h_over_tau_k}$$

is fulfilled. The second condition means that at initial time, wavelengths should be sufficiently far outside the Hubble scale; the first condition, which over-seeds the second one for the smallest wave numbers, means that the initial time should not be too close to recombination.

- Two parameters define the time at which the TCA is turned off for each wave number. This happens when one of the two conditions

$$(\tau_c/\tau_H) \geq \text{tight_coupling_trigger_tau_c_over_tau_h}$$

or

$$(\tau_c/\tau_k) \geq \text{tight_coupling_trigger_tau_c_over_tau_k}$$

is fulfilled. `CLASS` imposes that the TCA switching time should always be chosen after the initial time, which means that the four parameters above should satisfy simple inequalities.

- One parameter defines the time (common to all wave numbers) at which the source functions (leading to the computation of temperature and polarisation C_l 's) start being sampled and stored. This happens when the condition

$$(\tau_c/\tau_H) = \text{start_sources_at_tau_c_over_tau_h}$$

is satisfied. This time can eventually be chosen during the tight-coupling regime for the smallest wave numbers.

	no-tca	tca1	tca2	tca3
tight_coupling_trigger_tau_c_over_tau_h	$4.1 \cdot 10^{-4}$	$7 \cdot 10^{-3}$	$8 \cdot 10^{-3}$	$9 \cdot 10^{-3}$
tight_coupling_trigger_tau_c_over_tau_k	$6.1 \cdot 10^{-5}$	$3 \cdot 10^{-2}$	$5 \cdot 10^{-2}$	$8 \cdot 10^{-2}$

Table 1. Four settings for the precision parameters governing the time at which the TCA is switched off.




Figure 4. Impact of precision parameters governing the TCA switching time, for temperature (left) and E-polarisation (right). We show the power spectrum of the three settings `tca1`, `tca2` and `tca3` compared to the reference spectrum `no-tca` (see table 1 for precision parameter values).

In order to show the impact of these parameters, we take the set of precision parameters defined in the file `cl_permille.pre` of the CLASS public distribution, which corresponds to an accuracy of at least 0.1% on each temperature and polarisation C_l , and uses the `compromise_CLASS` scheme. We then vary the two trigger parameters mentioned above, as described in table 1. The first setting, called `no-tca`, corresponds to switching off the tight coupling approximation immediately after setting the initial conditions, so that no TCA is ever used. This leads to reference temperature/polarisation spectra with respect to which all the other results of this section are compared. The settings called `tca1`, `tca2` and `tca3` introduce from 0.02% to 0.08% of error with respect to the `no-tca` case, as illustrated in figure 4.

In figure 5, we stick to the precision setting `tca3` and compare the different TCA schemes. The `first_order_CAMB` and `first_order_CLASS` results are indistinguishable, confirming the fact that the approximation $c_s^2 \propto a^{-1}$ is sufficient in practice. Our second-order results and those derived from [12] are also in perfect agreement. As expected, the results from the `compromise_CLASS` scheme are essentially as good as the full second-order results, while the first-order results are roughly ten times less accurate. We also show on this plot the error produced by the `first_order_CLASS` scheme with `tca1` precision settings, which is similar to that produced by the `compromise_CLASS` scheme with `tca3` precision settings. Hence, in order to estimate the usefulness of going beyond the first order TCA, we can compare the performances of the code in these two cases.




Figure 5. Impact of TCA schemes on temperature (left) and E-polarisation (right). For the various TCA schemes discussed in section 2.5, we show the power spectrum with accuracy settings `tca3` compared to the reference spectrum `no-tca` (see table 1 for precision parameter values). We show a single curve for cases which are indistinguishable by eye, namely: `first_order_CAMB` and `first_order_CLASS`; and also, `second_order_CRS` and `second_order_CLASS`. The default scheme `compromise_CLASS` is also hardly distinguishable from the `second_order_CLASS`. The faint line shows for comparison the first-order results with accuracy `tca1`: the error is then comparable to `compromise_CLASS` with `tca3`.

precision	<code>no-tca</code>	<code>tca1</code>	<code>tca3</code>
TCA scheme	(irrelevant)	<code>first_order_class</code>	<code>compromise_class</code>
<code>rk</code>	1069s	19.4s	17.8s
<code>ndf15</code>	16s	14.9s	14.6s

Table 2. Execution time of the perturbation module, in seconds, with several TCA settings and with the two integrators (Runge-Kutta and stiff integrator `ndf15`). The last two columns lead to roughly the same level of accuracy.

In table 2, we compare running times in the `no-tca` case and in the previous two cases, using either the Runge-Kutta or `ndf15` integrator. The timings displayed here correspond to the number of seconds spent by our computer in the perturbation module of `CLASS`, in a non-parallel execution. The `ndf15` integrator is always better, by a huge amount in the `no-tca` case, or by 20 to 30% in the other cases. Being unaffected by the issue of integrating a stiff system, the `ndf15` integrator is not very sensitive to the choice of TCA scheme, with only a 3% speed up when using the compromise scheme instead of first-order schemes. The Runge-Kutta integrator is more sensitive, with a 9% speed-up for the compromise scheme.

We conclude that `CLASS` benefits much more from the implementation of our stiff integrator than from going beyond the first-order TCA. For some particular models, the user may wish to stick to the Runge-Kutta integrator, in which case the `compromise_CLASS` scheme leads to a sizable speed up.

3 Ultra-relativistic Fluid Approximation (UFA)

All massless neutrinos and ultra-relativistic relics can be treated as a single species, labeled as “ur” in CLASS. The code assumes that these species are fully decoupled. Hence, they just free-stream within a given gravitational potential, and can be followed with the collisionless Boltzmann equation expanded in harmonic space and integrated over momentum [10]. The solution can be formally written as the sum of spherical Bessel functions $j_l(k\tau)$ (exhibiting damped oscillations for $\tau > l/k$), plus a particular solution of the inhomogeneous equations sourced by metric fluctuations. The fact that one part of the solution has an analytical expression cannot be directly implemented in the code, because the total perturbations δ_{ur} , θ_{ur} and σ_{ur} back-react on the metric perturbations through Einstein equations, and affect the source terms in the Boltzmann equation. We will use this decomposition only as a guideline for deriving accurate approximation schemes.

3.1 Truncation of the Boltzmann hierarchy

Since the ur species couple only gravitationally to other species, we are only interested in tracking δ_{ur} , θ_{ur} and σ_{ur} . Higher multipoles must still be included since they couple with the lower ones, but in all efficient Boltzmann codes, the hierarchy is truncated at some low multipole value l_{max} . CMBFAST, CAMB, CMBEASY and CLASS all use the truncation scheme proposed in Ma & Bertschinger (eq. (51) of [10]) which is designed to minimize artificial reflection of power from l_{max} back to lower multipoles. Still, this truncation is not perfect, and a significant amount of unphysical reflection cannot be avoided for times beyond $\tau = l_{\text{max}}/k$. This implies that in order to compute an accurate CMB spectrum, l_{max} should be at least of the order of 30. The computation of the matter power spectrum $P(k)$ on small scales, up to some wavenumber k_{max} , is more problematic: one should further increase l_{max} proportionally to k_{max} in order to get converging results.

3.2 Sub-Hubble fluid approximation

The Ultra-relativistic Fluid Approximation (UFA) implemented in CLASS is based on the idea that for a given wavenumber, l_{max} should not necessarily be fixed throughout the whole time evolution. The code considers two regimes: wavelengths larger or comparable to the Hubble radius, and wavelengths much smaller than the Hubble radius. The transition between the two regime occurs for each k when the product $k\tau$ (equal to τ/τ_k and coinciding with τ_H/τ_k during radiation domination) reaches some threshold value that we call here $(k\tau)_{\text{ufa}}$. Typically, $(k\tau)_{\text{ufa}}$ is chosen in the range from 10 to 50, depending on the required precision. The full name of this parameter in the code is `ur_fluid_trigger_tau_over_tau_k`. In the first regime $k\tau \leq (k\tau)_{\text{ufa}}$, the Boltzmann hierarchy can be truncated at some l_{max} which can be chosen to be rather small: it is enough to take to $l_{\text{max}} \sim (k\tau)_{\text{ufa}}$, since multipoles with $l > k\tau$ are negligible (according to the spherical Bessel function approximation). In the second regime $k\tau \geq (k\tau)_{\text{ufa}}$, the code still follows the three variables δ_{ur} , θ_{ur} and σ_{ur} , which are sourced by metric perturbations. But multipoles in the range $2 < l \ll k\tau$ are suppressed, leading to an effective decoupling between the first three multipoles and the highest ones. Hence it is natural to lower l_{max} down to two in this regime. Ultra-relativistic neutrinos are then described by a reduced system of equations for δ_{ur} , θ_{ur} and σ_{ur} , i.e. by fluid equations (of course, this fluid is not assumed to be perfect, since it has anisotropic pressure). In summary, the UFA approximation consists in lowering l_{max} from a value close to $(k\tau)_{\text{ufa}}$

down to $l_{\text{max}} = 2$ deep inside the Hubble radius, at the time when $k\tau = (k\tau)_{\text{ufa}}$. Such a scheme offers many advantages:

1. When computing the matter power spectrum, the number of **ur** equations to integrate before the approximation is switched on does not need to be scaled linearly with the highest wave number k_{max} .
2. The number of **ur** equations reduces to $(l_{\text{max}} + 1) = 3$ in the whole region of (k, τ) space fulfilling the condition $k\tau > (k\tau)_{\text{ufa}}$; this is precisely the region in which the computation would be time-consuming, since **ur** perturbations oscillate inside the Hubble radius.
3. The UFA completely avoids the issue of power reflecting at some large l_{max} , which would otherwise affect the evolution of low multipoles periodically due to some spurious wave travelling back and forth between $l = l_{\text{max}}$ and $l = 0$. (This behaviour is clearly seen in figure 6.)

The fluid approximation could in principle be used until present time, but the code allows a more aggressive approximation, the Radiation Streaming Approximation, to take over from the UFA after photon decoupling. This new approximation is discussed in the next section. Hence, the UFA is essentially a way to save computing time during radiation domination and at the beginning of matter domination.

3.3 Fluid equations

We need to find a closed system for the evolution of δ_{ur} , θ_{ur} and σ_{ur} , valid deep inside the Hubble radius. The full system of equations in the synchronous gauge (eq. (49) in [10]) reads:

$$\delta_{\text{ur}}' = -\frac{4}{3}\theta_{\text{ur}} - \frac{2}{3}h', \quad (3.1a)$$

$$\theta_{\text{ur}}' = k^2 \left(\frac{1}{4}\delta_{\text{ur}} - \sigma_{\text{ur}} \right), \quad (3.1b)$$

$$2\sigma_{\text{ur}}' = \frac{8}{15}\theta_{\text{ur}} - \frac{3}{5}kF_{\text{ur}3} + \frac{4}{15}(h' + 6\eta'), \quad (3.1c)$$

$$F_{\text{ur}l}' = \frac{k}{2l+1} [lF_{\text{ur}(l-1)} - (l+1)F_{\text{ur}(l+1)}]. \quad (3.1d)$$

In appendix B, we use the formal solution of these equations in order to derive an exact integral relation between σ_{ur}' , σ_{ur} , θ_{ur} and metric perturbations. We then find an approximate but more practical form of this relation valid inside the Hubble radius, at leading order in an expansion in metric perturbation derivatives ($h^{(n)}/k^{n-1}$, $\eta^{(n)}/k^{n-1}$) and in powers of $(k\tau)^{-1}$:

$$\sigma_{\text{ur}}' = -\frac{3}{\tau}\sigma_{\text{ur}} + \frac{2}{3}\theta_{\text{ur}} + \frac{1}{3}h'. \quad (3.2)$$

Since metric perturbation only evolve over a Hubble time scale inside the Hubble radius, we expect this expansion to converge, and we will see below that the above relation is indeed accurate enough for our purpose. In the default version of CLASS, this equation is used for closing the system of equations when the UFA is switched on. This method corresponds to the setting `ufa_method = ufa_class` in the code's precision parameter structure.

3.4 Alternative schemes

Some nearly equivalent schemes can be justified in slightly different ways. Truncating the Boltzmann equations at $l_{\max} = 2$ with the usual truncation scheme of Ma & Bertschinger gives:

$$\sigma_{\text{ur}}' = -\frac{3}{\tau}\sigma_{\text{ur}} + \frac{2}{3}\theta_{\text{ur}} + \frac{1}{3}(h' + 6\eta'). \quad (3.3)$$

This truncation scheme is based entirely on the assumption that $F_{\text{ur } l}(k, \tau) \propto j_l(k\tau)$. So, the reason for the difference between eq. (3.2) and (3.3) is that (3.2) is based on the full formal solution, including the leading order contribution to the part sourced by the metric, while (3.3) is based only on the solution of the homogeneous equation. Equation (3.3) is used when the user switches to `ufa_method = ufa_mb`, and amounts to adding an extra term in η' .

Finally, in a more general context, Hu [13] introduced a set of equations modeling a cosmological viscous fluid, and suggested that this fluid could approximate the evolution of free-streaming neutrinos with the parameter choice $(w, c_s^2, c_{\text{vis}}^2) = (1/3, 1/3, 1/3)$. In this limit, Hu's fluid equations are identical to our UFA equations except for the shear derivative:

$$\sigma_{\text{ur}}' = -3\frac{a'}{a}\sigma_{\text{ur}} + \frac{2}{3}\theta_{\text{ur}} + \frac{1}{3}(h' + 6\eta'). \quad (3.4)$$

The coefficient $-3\frac{a'}{a}$ reduces to $-3/\tau$ deep inside the radiation dominated regime, but becomes different around the time of equality. This equation is used when the user switches to `ufa_method = ufa_hu`. Below we will compare the performances of equations (3.2), (3.3), (3.4) and show that the first one is slightly more precise (as expected from the rigorous mathematical proof of appendix B).

Finally, when the user selects `ufa_method = ufa_none`, no UFA scheme is employed, and the truncation multipole l_{\max} remains the same throughout the evolution.

3.5 Comparison at the level of perturbations

In figure 6, we compare the evolution of δ_{ur} and σ_{ur} for a given mode, obtained either by solving the full Boltzmann equation up to a very high $l_{\max} \sim 3000$, or with $l_{\max} = 46$ with/without the default UFA. In absence of approximation, one can see some spurious evolution appearing periodically (here, around $k\tau \sim 100$ and $k\tau \sim 200$): this corresponds to the propagation of an unphysical wave between the multipole boundaries $l = l_{\max}$ and $l = 0$. Using the default UFA scheme `ufa_class`, we reproduce accurately the phase, the amplitude and, to a lesser extent, the zero-point of the oscillations. For the clarity of the figure, we do not show the results from alternative approximation schemes. We checked that the `ufa_mb` scheme also reproduces the correct phase and amplitude, but introduces a larger error in the zero-point of oscillations. Finally, the `ufa_hu` scheme reproduces the phase, but not the correct amplitude of the oscillations.

3.6 Comparison at the level of CMB/matter power spectrum

In figure 7, we show the impact of the UFA on the CMB and matter power spectrum. We take the precision parameters of the file `cl_permille.pre`, and play with the values of l_{\max} and $(k\tau)_{\text{ufa}}$, called `l_max_ur` and `ur_fluid_trigger_tau_over_tau_k` in the code. We first compute some reference spectra with $l_{\max} = 3000$ (to remove any truncation effect), and such a large value of $(k\tau)_{\text{ufa}}$ that the UFA is never used. All other results from this section are compared to these spectra. We then fix both l_{\max} and $(k\tau)_{\text{ufa}}$ to 18 and vary




Figure 6. Evolution of δ_{ur} (left) and σ_{ur} (right) for the mode $k = 10^{-1} h \text{Mpc}^{-1}$ over the range $10 < k\tau < 200$, i.e. well inside the Hubble radius and throughout the radiation dominated stage. In the “exact” case, the full Boltzmann equation for ur is truncated at $l_{\text{max}} \sim 3000$, with no impact of the truncation on the result. In the “no approx.” case, the truncation is performed at $l_{\text{max}} = 46$. In the “CLASS” case, we use the UFA scheme `ufa_class` and set $l_{\text{max}} = 46$ as long as $k\tau \leq (k\tau)_{\text{ufa}} = 50$, or $l_{\text{max}} = 2$ afterward.

UFA scheme (fixed l_{max} and $(k\tau)_{\text{ufa}}$)	<code>ufa_none</code>	<code>ufa_class</code>
<code>rk</code>	29.7s	27.0s
<code>ndf15</code>	16.7s	15.2s

Table 3. Execution time of the perturbation module in seconds, with the precision parameters of the file `cl_permille.pre`, plus $l_{\text{max}} = (k\tau)_{\text{ufa}} = 18$. Using an ultra-relativistic fluid approximation leads simultaneously to a 10% faster execution and to more accurate results.

only the `ur_fluid_approximation` parameter. We show the error induced by each UFA scheme for the temperature and matter power spectrum in figure 7 (results for temperature and polarisation are very similar). The results from the `ufa_none` case are very unstable and depend a lot on the choice of l_{max} and $(k\tau)_{\text{ufa}}$ values. With the present choice, they correspond to a twice larger error in the CMB spectra than in any UFA scheme; for slightly different choices they would also induce a larger error in the matter power spectrum. The three UFA schemes, which do not have such instabilities, are nearly as good as each other for CMB spectra, while for the matter power spectrum the `ufa_class` scheme is one order of magnitude better. Table 3 shows that the UFA approximation allows for a 10% speed up, while being more accurate for a fixed l_{max} .

4 Radiation Streaming Approximation (RSA)

After their respective decoupling time, the photons gradually free-stream like neutrinos (except around the time of reionisation at which their coupling to baryons is enhanced). In principle, it would be possible to look for a fluid approximation for photons, like we did for neutrinos in the previous section. However we can go further than that, since during this




Figure 7. Power spectra (for temperature CMB anisotropies and for matter) obtained with various implementations of the UFA, compared to those obtained in a reference run. In the four cases displayed here, we use $l_{\max} = 18$ as long as $k\tau \leq (k\tau)_{\text{ufa}} = 18$. The curves labeled “MB”, “HU”, “CLASS” correspond to the three possible implementation of the UFA discussed in the text; the last curve uses no UFA and a standard truncation scheme at $l_{\max} = 18$ until the time at which the next approximation RSA takes over (see section 4): in this case the code is at the same time a bit less precise and 10% slower. We do not show the results for the polarisation spectrum C_l^{EE} , that look very similar to those for C_l^{TT} .

period the universe is dominated by matter and eventually Λ /Dark Energy: in this case photons and massless neutrinos almost behave like test-particles in an external gravitational field, and we do not need to catch their evolution with high accuracy (which was not the case for the `ur` species during radiation domination).

Like for massless neutrinos, in all efficient Boltzmann codes, the Boltzmann equation for photons is truncated at some low multipole value l_{\max} using the truncation scheme proposed in Ma & Bertschinger (eqs. (65) in [10]). If l_{\max} is not large enough, the spurious reflection of power induced by the truncation propagates to the final results, because radiation perturbations still play a small role during the free-streaming regime. More precisely:

- the photon density fluctuation, velocity and shear perturbations appear in Einstein equations;
- the photon density fluctuation and shear appear in the temperature/polarisation source functions;
- the photon velocity appears in the evolution equations of baryons (since the baryon-photon coupling is not negligible during reionisation).

In order to avoid propagating such an error, one can either increase l_{\max} , or find a way to infer the photon density, velocity and shear from some analytic Radiation Streaming Approximation (RSA), in which case the integration of Boltzmann equations can be stopped soon after photon decoupling. Mathematically, this analytic approximation should coincide with the particular non-oscillatory solution of the inhomogeneous Boltzmann equations. Once

the damped oscillations accounted for by spherical Bessel functions becomes negligible (i.e. when $k\tau \gg l$), the analytic approximation will coincide with the true solution. Before that time, it will provide a correct approximation to the true quantities averaged over a few oscillations. In summary, the RSA has a double goal: to avoid unphysical oscillations created by the Boltzmann truncation, and to avoid wasting time in integrating the Boltzmann equations over many such oscillations.

The RSA does not need to be accurate at very late time (end of matter domination, Λ /Dark Energy domination), since by then radiation fluctuations are always completely negligible with respect to matter fluctuations. However, it should be reasonably accurate soon after photon decoupling, i.e. during matter domination, when the energy density in radiation, $\Omega_r \equiv \Omega_\gamma + \Omega_{\text{ur}}$, is smaller than one but not yet much smaller. The advantage of a better approximation is two-fold. First, it can be switched on earlier. Second, before switching on the approximation, we can use a smaller value of l_{max} , since high multipoles will not have time to grow.

The same treatment can be applied to ultra-relativistic species, which are identical to photons in this regime, except for the fact that they remain collisionless during reionisation. When the RSA is turned on for photons, it is better to follow ultra-relativistic species in the same way as photons, rather than with the fluid formalism described in section 3. The RSA then removes three more differential equations, and cures the fact that the UFA turns out to be inaccurate at late time. Hence, the default version of CLASS treats ultra-relativistic species first with exact equations, then with the UFA (inside the Hubble radius and until photon decoupling), and finally with the RSA (inside the Hubble radius and after photon decoupling).

An expression for the RSA was discussed in the Newtonian gauge by Doran [11]. Soon after, a somewhat simpler RSA (neglecting reionisation) was also introduced in the CAMB code, which uses the synchronous gauge. Here, we will derive an approximation comparable to that of Doran [11], but valid in the synchronous gauge.

4.1 Relativistic relics (massless neutrinos)

We start with the simplest case, that of ultra-relativistic species ur . We combine the first two equations of (3.1) into

$$\delta_{\text{ur}}'' + \frac{k^2}{3}\delta_{\text{ur}} = -\frac{2}{3}h'' + \frac{4}{3}k^2\sigma_{\text{ur}}. \quad (4.1)$$

Inside the Hubble scale (i.e. when $k\tau \gg 1$) we can assume in first approximation that $|\sigma_{\text{ur}}| \ll |\delta_{\text{ur}}|$ and neglect the shear in the RSA. Also, since we are looking for a smooth (non-oscillatory) particular solution of this inhomogeneous equation, we can assume that $|\delta_{\text{ur}}''| \ll k^2|\delta_{\text{ur}}|$. We conclude that the RSA for δ_{ur} is simply

$$\delta_{\text{ur}} = -\frac{2}{k^2}h''. \quad (4.2)$$

Note that in the synchronous gauge, h' coincides with $-2\delta'_{\text{cdm}}$, where δ_{cdm} is the cold dark matter density contrast. Deep inside the matter-dominated regime, $\delta_{\text{cdm}} \propto a \propto \tau^2$, so h' is linear in τ , and h'' is a constant. The RSA for δ_{ur} is therefore nearly static. Concerning θ_{ur} , its value in the RSA is given by the exact energy-conservation equation $\delta'_{\text{ur}} = -\frac{4}{3}\theta_{\text{ur}} - \frac{2}{3}h' = 0$, namely

$$\theta_{\text{ur}} = -\frac{1}{2}h'. \quad (4.3)$$

In practice, we must extract h' and h'' from the Einstein equations in the synchronous gauge, that read

$$2k^2\eta - \frac{a'}{a}h' = 8\pi\mathcal{G}a^2\delta\rho_{\text{tot}}, \quad (4.4a)$$

$$2k^2\eta' = 8\pi\mathcal{G}a^2[(\bar{\rho} + \bar{p})\theta]_{\text{tot}}, \quad (4.4b)$$

$$h'' + 2\frac{a'}{a}h' - 2k^2\eta = -8\pi\mathcal{G}a^2\delta p_{\text{tot}}. \quad (4.4c)$$

Here we do not need the fourth equation sourced by the shear. The difficulty comes from the fact that in order to infer δ_{ur} we should compute h'' , and for doing that we need to combine the first and third equation, i.e. we need to know $\delta\rho_{\text{tot}}$, which depends itself on δ_{ur} . Fortunately, we can notice that if we omit δ_{ur} in the computation of $\delta\rho_{\text{tot}}$, we make a tiny error, since during matter domination $|\delta\rho_{\text{ur}}| \ll |\delta\rho_{\text{cdm}}|$. Hence, it is good enough to evaluate the first Einstein equation with δ_{ur} set to zero. The same is not true for the second equation, since the synchronous gauge is comoving with cdm, so one has $\theta_{\text{cdm}} = 0$ by construction. As a result, neglecting θ_{ur} in the computation of θ_{tot} and η' leads to a significant inaccuracy in the solution for η . Hence, we introduce the following scheme:

1. We compute $\delta\rho_{\text{tot}}$ assuming $\delta_{\text{ur}} = 0$, and obtain $2k^2\eta - \frac{a'}{a}h'$ from the first Einstein equation.
2. Using the fact that during matter domination $|\delta p_{\text{tot}}| \ll |\delta\rho_{\text{tot}}|$, we notice that

$$\left| h'' + 2\frac{a'}{a}h' - 2k^2\eta \right| \ll \left| \frac{a'}{a}h' - 2k^2\eta \right|, \quad (4.5)$$

and hence to very good approximation

$$h'' = -2\frac{a'}{a}h' + 2k^2\eta. \quad (4.6)$$

We then infer the following RSA for δ_{ur} from eq. (4.2):

$$\delta_{\text{ur}} = \frac{4}{k^2} \left(\frac{a'}{a}h' - k^2\eta \right). \quad (4.7)$$

This formula is practical since η is one of the variables that we integrate over time, and h' has been inferred in the previous step.

3. We impose the free-streaming solution for θ_{ur} (eq. (4.3)) and set $\sigma_{\text{ur}} = 0$.
4. We use the remaining Einstein, continuity and Euler equations to evolve the system.

4.2 Photons

For photons, the solution is a bit more complicated since the baryon-photon coupling cannot be neglected during reionisation. We then need to find the particular non-oscillatory solution of (cf. (2.4))

$$\delta''_{\gamma} + \frac{k^2}{3}\delta_{\gamma} = -\frac{2}{3}h'' + \frac{4}{3}k^2\sigma_{\gamma} - \frac{4}{3\tau_c}(\theta_b - \theta_{\gamma}). \quad (4.8)$$

Once again we will neglect the shear and search for a particular solution slowly varying with time ($|\delta''_{\gamma}| \ll |k^2\delta_{\gamma}|$). In order to deal with the coupling term, we expand the solution in

powers of τ_c^{-1} . The zeroth-order solution is exactly similar to that of massless neutrinos, eqs. (4.7), (4.3). The first-order solution should satisfy

$$\frac{k^2}{3}\delta_\gamma = -\frac{2}{3}h'' - \frac{4}{3\tau_c}\left(\theta_b + \frac{1}{2}h'\right), \quad (4.9)$$

in which h'' can be replaced using eq. (4.6). This approximation turns out to work out very well, unlike the zeroth-order solution. The velocity is then given by the exact energy-conservation equation

$$\theta_\gamma = -\frac{1}{2}h' - \frac{3}{4}\delta'_\gamma. \quad (4.10)$$

We take the derivative of the previous result for δ_γ , assuming that h'' is time-independent, and using once more eq. (4.6). We obtain:

$$\theta_\gamma = -\frac{1}{2}h' + \frac{3}{k^2\tau_c}\left[-\frac{\tau'_c}{\tau_c}\left(\theta_b + \frac{1}{2}h'\right) + \left(\theta'_b + \frac{1}{2}h''\right)\right], \quad (4.11)$$

in which h'' can be replaced using eq. (4.6). However the exact expression of θ'_b depends again on θ_γ . Like before, we use a perturbative scheme in τ_c^{-1} and replace θ'_b above by its expression at first-order in τ_c^{-1} .

4.3 Summary of RSA equations

In summary, the RSA consists in neglecting δ_γ and δ_{ur} in the evolution of $\delta\rho_{\text{tot}}$ in the first Einstein equation, which allows us to obtain h' , and then in imposing

$$\delta_\gamma = \frac{4}{k^2}\left(\frac{a'}{a}h' - k^2\eta\right) + \frac{4}{k^2\tau_c}\left(\theta_b + \frac{1}{2}h'\right), \quad (4.12a)$$

$$\theta_\gamma = -\frac{1}{2}h' + \frac{3}{k^2\tau_c}\left[-\frac{\tau'_c}{\tau_c}\left(\theta_b + \frac{1}{2}h'\right) + \left(-\frac{a'}{a}\theta_b + c_b^2k^2\delta_b - \frac{a'}{a}h' + k^2\eta\right)\right], \quad (4.12b)$$

$$\sigma_\gamma = 0, \quad (4.12c)$$

$$\delta_{\text{ur}} = \frac{4}{k^2}\left(\frac{a'}{a}h' - k^2\eta\right), \quad (4.12d)$$

$$\theta_{\text{ur}} = -\frac{1}{2}h', \quad (4.12e)$$

$$\sigma_{\text{ur}} = 0. \quad (4.12f)$$

This scheme is set to be the default one in CLASS, as long as the precision variable `radiation_streaming_approximation` remains set to `rsa_MD_with_reio`. For comparison, some cruder schemes can be used: if the same variable is set to `rsa_MD`, the code will use the above expressions at zero order in τ_c^{-1} (i.e, with $\delta_\gamma = \delta_{\text{ur}}$ and $\theta_\gamma = \theta_{\text{ur}}$). If it is set to `rsa_none`, the radiation perturbations are just set to zero.

4.4 Comparison at the level of perturbations

In figure 8, we show the evolution of δ_γ , θ_γ , θ_b and η between recombination and present time, for a particular wavenumber $k = 0.1 \text{ Mpc}^{-1}$. We compare two RSA schemes with the exact evolution obtained by integrating all multipoles at all times. We always keep the UFA approximation of section 3 turned off. We see that immediately after switching on the RSA, our approximation for δ_γ (and also for δ_{ur} , which is not shown) matches accurately the exact




Figure 8. Evolution of the quantities δ_γ (top left), θ_γ (top right), θ_b (bottom left), η (bottom right) for the mode $k = 0.1 \text{ Mpc}^{-1}$, as a function of conformal time (in Mpc), between a time chosen soon after photon decoupling (or slightly before reionisation in the θ_b plot) and today. The red curves show the result of an exact integration with no truncation or approximation. In the blue curves, the RSA is turned on around $\eta = 470 \text{ Mpc}$, just before the solutions become unphysical due to the Boltzmann equation truncation at $l_{\text{max_g}}=12$, $l_{\text{max_pol_g}}=12$, $l_{\text{max_ur}}=28$. For comparison, in green, we show the result for δ_γ when the terms in τ_c^{-1} are neglected in the RSA, and those for θ_b and η when the radiation multipoles are all set to zero instead of using a free-streaming solution.

evolution averaged over a few oscillations. This would not be the case with several simpler RSA schemes which assume full matter domination and an exact linear growth of $h'(\tau)$. In the models used for the figures, reionisation takes place at $z_* = 10$ and $\tau = 4458 \text{ Mpc}$. It induces a clear feature in δ_γ and θ_b (having impact on η) which is well captured by the terms proportional to τ_c^{-1} in the full `rsa_MD_with_reio` scheme.

4.5 Comparison at the level of temperature/polarisation spectra

The precision parameters governing the evolution of perturbations in the late universe are:

- two parameters defining the time at which the RSA is switched on. For each wavenumber, we stop evolving photon and ur perturbations when the two conditions

$$k\tau = \tau/\tau_k \geq \text{radiation_streaming_trigger_tau_over_tau_k}$$

and

$$\tau_c/\tau \geq \text{radiation_streaming_trigger_tau_c_over_tau}$$

are satisfied (i.e., photons are sufficiently decoupled, and the wavelength is sufficiently deep inside the sub-Hubble regime). In principle, it would be possible to switch on the RSA at different times for photons and ur species, but for simplicity we did not consider this option.

- `l_max_g`, `l_max_pol_g` and `l_max_ur` define the number of photon temperature, photon polarisation and ur multipoles which are integrated until the RSA is switched on.

precision setting:	reference	no-rsa	rsa
<code>l_max_g</code>	3000	18	18
<code>l_max_pol_g</code>	3000	18	18
<code>l_max_ur</code>	3000	50	50
<code>radiation_streaming_trigger_tau_over_tau_k</code>	∞	∞	100
<code>radiation_streaming_trigger_tau_c_over_tau</code>	∞	∞	2
<code>ur_fluid_trigger_tau_over_tau_k</code>	∞	∞	∞

Table 4. Three settings for the parameters governing the Boltzmann truncation and RSA. All other parameters are fixed with the file `cl_permille.pre` of the public CLASS distribution: in particular, `radiation_streaming_approximation` is set to `rsa_MD_with_reio`. The reference run never uses the UFA and RSA, and cannot be affected by the Boltzmann truncation. The second and third settings share the same truncation multipoles, and differ only by using the RSA or not.




Figure 9. Temperature and matter spectra for the runs `rsa` and `no-rsa`, normalized by reference spectra. Corresponding precision parameter settings are described in table 4. Using the RSA with such settings leads to equally accurate CMB spectra, and to a very small error on the matter power spectrum. The running time is however reduced considerably.

In order to compute some reference spectra to be used throughout this section, we fix the precision parameters according to the file `cl_permille.pre`, increase `l_max_g`, `l_max_pol_g` and `l_max_ur` to 3000, and choose such large values of the trigger parameters that the UFA and RSA are never employed. We wish to compare these reference spectra with those from runs with/without the RSA, in which `l_max_g`, `l_max_pol_g` and `l_max_ur` are kept fixed to reasonable values. We choose `l_max_g` and `l_max_pol_g` to be equal to 18. Since we do not want to use the UFA approximation in this comparison (in order to focus only on the impact of the RSA), we fix `l_max_ur` to a larger value, namely 50. This setting, called `no-rsa` in table 4, leads to a 0.01% error both in the temperature multipoles and in the matter power spectrum for $k \leq 1 h \text{ Mpc}^{-1}$, as shown in figure 9. Finally, in the run called `rsa`, we switch on the default RSA scheme, with the trigger values specified in table 4. With such settings, the error in the temperature (and also polarisation) multipoles remains as small as without the

Runs with/without a RSA:	no-rsa	rsa
rk	55.8s	33.6s
ndf15	84.4s	14.5s

Table 5. Execution time of the perturbation module in seconds, with the precision parameters of the file `cl_permille.pre`, modified as described in table 4.

RSA, while the error in the matter power spectrum grows moderately to 0.04% (see figure 9). However, the running time is reduced considerably, as shown in table 5: with the Runge-Kutta integrator, the RSA leads to a 66% speed up. Note that the `ndf15` remains much better than the Runge-Kutta integrator when the RSA is employed (by a factor 2) while it experiences difficulties in following oscillatory solutions in absence of a RSA. However, the combination of our RSA scheme and `ndf15` integrator leads to very nice performances (speed-up by a factor 4 with respect to Runge-Kutta without any RSA).

In CAMB, the RSA is somewhat cruder, since it neglects reionisation and uses an explicit cosmology-dependent relation giving h' an h'' in terms of $\bar{\rho}_b$, $\bar{\rho}_{\text{cdm}}$, δ_b , δ_{cdm} and k , valid only deep in the matter-dominated regime. We present in a companion paper the comparison between the matter power spectrum $P(k)$ computed by CAMB and CLASS. In order to get an accurate $P(k)$ with CAMB, one is forced to deactivate the RSA approximation (called `late_rad_truncation`), precisely for the above reasons. Our scheme does not lead to a significant error on the $P(k)$, and is more model-independent: it involves only metric perturbations and does not need to be modified in the presence of other components playing a role during matter domination (e.g. with warm dark matter or early dark energy).

5 Conclusions

In figure 10, we summarise the different approximations used by the code in the (k, τ) plane when computing the C_l 's up to 3000 and the $P(k)$ up to $1 h \text{ Mpc}^{-1}$. The figure corresponds to the precision settings of the file `cl_3permille.pre`. Exact equations are used only in the band corresponding to Hubble crossing for each mode, as well as in the super-Hubble region with non-tightly-coupled photons, in which all quantities evolve very slowly and integration is very fast.

With these approximations, for ΛCDM , the perturbation module only spends a significant time in the region corresponding to Hubble crossing for each mode, and to the sub-Hubble evolution before photon decoupling. At early times, stiff equations are avoided thanks to the Tight-Coupling Approximation. Well-inside the Hubble radius and until photon decoupling, the different modes oscillate and integration is time-consuming: however, the number of equations is kept small (of the order of 30 in total) thanks to the UFA. After photon decoupling, the code only needs to integrate over 4 equations with very smooth solutions, and the time spent by the code in the RSA region is negligible.

Other approximation schemes can be introduced in more general cosmological models. The case of massive neutrinos and non-cold dark matter relics is discussed in a companion paper. More exotic cases may require further approximations which can be introduced and discussed case-by-case (CLASS is coded in such way that introducing a new approximation is as structured, codified and simple as introducing new species [6]). However, the fact that CLASS uses an original stiff integrator means that for several purposes (as for the generalisation of TCA), new approximation schemes are not even strictly necessary.




Figure 10. Summary of regions in (k, τ) space where the various approximations are used. The precision settings are taken from the `cl_3permille.pre` precision file which ensures a 0.3% precision on the C_l 's till $l = 3000$. The full set of exact equations is used only in the white region.

Acknowledgments

This work was supported in part by the Swiss Science Foundation (D.B.). Running with high-precision settings is only practical on machines with many cores and large memory: we wish to thank M. Shaposhnikov for providing us with a brand new 48-core PC at EPFL, and the Institut d'Astrophysique de Paris for giving us access to the horizon9 machine. Finally, we would like to thank A. Lewis for many useful comments.

A Stiff integrator

The standard numerical method for solving Ordinary Differential Equations (ODEs) is to use an adaptive step size Runge-Kutta solver. While this method is fast and accurate in simple cases, it may fail completely (or take a very large number of steps) when the problem is stiff. Stiffness occurs when at least two time scales of evolution in the problem differ substantially. A well known example is the Boltzmann equation in cosmology. If a distribution is kept in equilibrium by the (rapid) interaction with a background species, and we are interested in the evolution of the distribution on cosmological time-scales, a Runge-Kutta solver will oscillate

around the equilibrium solution by means of very small time steps related to the time-scale of the interaction.

This problem exists, for instance, in the early universe where baryons and photons are strongly coupled. One solution is to substitute the equilibrium solution into the equations and thereby reducing the system of equations and removing the stiffness: this is the basic principle of the Tight Coupling Approximation (TCA) discussed in section 2. This approximation removes the stiffness of the Boltzmann equation in the vanilla scenario.

However, stiffness may easily be reintroduced by trying to incorporate new physics into the code. The proper TCA approximation should then be derived and implemented in the code, which would require a great deal of familiarity with the code from the user. In CLASS, however, the user can just take advantage of the implemented stiff solver.

The `ndf15` algorithm is a variable order (1-5) adaptive step size solver based on the Numerical Differentiation Formulas of order 1 to 5. The step size is adaptive but quasi-constant, meaning that the formulas used are based on a fixed step size. Each time the step size changes, the code will update the backward differences by interpolation to reflect this new step size. The algorithm is described in [17].

Whereas Runge-Kutta methods are explicit, meaning that the next step can be computed directly from the previous step by elementary operations, `ndf15` is a fully implicit method. This means that at each time step, we must solve a system of algebraic, possibly non-linear equations, which is accomplished by Newton iteration. This requires a numerical computation of the Jacobian and the solution of systems of linear equations. As it is evident, all this can make each step very expensive, so a number of strategies must be implemented:

- Reusing Jacobians:

The Jacobian usually changes more slowly than the solution itself, so an attempt is made to reuse the Jacobian - it will only be recomputed if Newton iteration is too slow.

- LU decomposition:

The same linear system must be solved repeatedly with different right hand sides, so we should of course store an LU decomposition. Because the system of equations can be large, $O(100)$, we need sparse matrix methods for this.

- Backward Interpolation:

Since the method stores a matrix of backward differences, it is fast to infer the value and the derivatives of the solution at points before the current point by interpolation. We only need the values of some of the components to calculate the source functions, so we do not need to interpolate the rest of them.

When the number of equations is larger than about 10, it is advantageous to use sparse matrix methods, and if the system is somewhat larger, the difference in execution time can differ by orders of magnitude. Usually sparse matrix methods are developed in order to save both time and memory, but for our purposes only execution time matters. We created a small sparse matrix package for our purpose based on [18]. Some important features of this package are:

- Column Pre-ordering:

The matrices which appear are of the general form $I - cJ$ where I is the identity matrix, J is the Jacobian and c is some constant. Since the Jacobian is close to being

structurally symmetric (if y_i couples to y_j , it will often be the case that y_j couples to y_i), it is advantageous to use the Approximate Minimum Degree column ordering of the matrix $C = J^T + J$. Using this pre-ordering reduces the number of non-zero elements in the corresponding L and U factors by a factor of a few, which leads to a sizable reduction in the time needed to solve the linear systems.

- LU re-factorisation:

If we have already calculated a LU factorization for some Jacobian, which is structurally identical to the current Jacobian, we can use information saved during the factorisation of the former to factorise the new matrix in a fraction of the time. Specifically, we store the pivot ordering and the *reach* of all the sparse right hand sides used in forming the LU-decomposition.

- Fast Jacobian Calculation:

If the same pattern of the Jacobian is found repeatedly, we can use this pattern to speed up the calculation. Taking advantage of the sparsity of the Jacobian, we can group the columns together and form the Jacobian using only a fraction of the usual n function evaluations, n being the number of equations.

The CLASS user can choose to use the Runge Kutta or `ndf15` integrator by switching the precision parameter `evolver` to either `rk` (=0) or `ndf15` (=1, default setting). To illustrate the power of `ndf15` in stiff situations, `ndf15` was less than 10% slower when the TCA was turned off as early as in the reference run used in figure 2. As a comparison, the standard Runge-Kutta integrator was more than 10000 times slower than `ndf15` for the same model. This particular run represents an extreme case, but throughout this work we have presented various examples in which the `ndf15` performances are very good.

B Derivation of fluid equations for ultra-relativistic relics

The goal of this appendix is to establish the validity of the approximate shear derivative equation (3.2), which allows to treat collisionless species as an imperfect fluid governed by may less equations. In the future, our results could be used for computing higher order terms in (3.2), or more generally for understanding various properties of the linear perturbations of ultra-relativistic species. A discussion similar in spirit was presented in [19] for massive neutrinos, although the goal of that paper was to introduce a sharp truncation at $l = 3$, while we are searching for a truncation scheme that would take into account the transfer of power to higher l 's.

B.1 Formal solution

Sticking to the notations of Ma & Bertschinger, the perturbations of ultra-relativistic species is described by a function $F(\mathbf{k}, \mu, \tau)$ obeying to the collisionless Boltzmann equation

$$\partial_\tau F(\mathbf{k}, \mu, \tau) + ik\mu F(\mathbf{k}, \mu, \tau) = S(\mathbf{k}, \mu, \tau), \quad (\text{B.1})$$

where S stands for the gravitational source terms. The most general formal solution can be written as

$$F(\mathbf{k}, \mu, \tau) = F^0(\mathbf{k}, \mu)e^{-ik\mu\tau} + \int_0^\tau e^{-ik\mu(\tau-\tilde{\tau})} S(\mathbf{k}, \mu, \tilde{\tau}) d\tilde{\tau}. \quad (\text{B.2})$$

The initial function $F^0(\mathbf{k}, \mu)$ depends on the considered gauge and type of initial conditions. For instance, in the synchronous gauge, $F^0 = 0$ for the growing ADiabatic (AD), Baryon Isocurvature (BI) and Cold Dark Isocurvature (CDI) modes; F^0 has a non-zero monopole term for Neutrino Isocurvature Density (NID) initial conditions; and a non-zero dipole term for Neutrino isocurvature Velocity (NIV). These statements can be checked from ref. [20]. Hence, in all cases, we can write

$$F^0(\mathbf{k}, \mu) = C_{\text{NID}}(\mathbf{k}) - ik\mu C_{\text{NIV}}(\mathbf{k}). \quad (\text{B.3})$$

More fundamentally, the fact that F^0 only has monopole and dipole contributions can be justified by the fact that neutrinos were initially in thermal equilibrium, forming a fluid with no anisotropic pressure or higher momenta. By causality, this remains true at any time after decoupling on super-Hubble scales. So, we are sure that the above form of F^0 is completely universal. In the synchronous gauge, the source reads

$$S = -\frac{2}{3}h' - \frac{4}{3}(h' + 6\eta')P_2(\mu). \quad (\text{B.4})$$

Thanks to a few integration by part, it is possible to absorb the μ dependence, in order to be able to write this solution in Legendre space. The result is

$$\begin{aligned} F(\mathbf{k}, \mu, \tau) &= F^0(\mathbf{k}, \mu)e^{-ik\mu\tau} \\ &+ \frac{2}{k^2} \int_0^\tau e^{ik\mu(\tilde{\tau}-\tau)} \{2k^2\eta'(\tilde{\tau}) + h'''(\tilde{\tau}) + 6\eta'''(\tilde{\tau})\} d\tilde{\tau} \\ &- \frac{2}{k^2} \{h'' + 6\eta'' - ik\mu(h' + 6\eta')\} \\ &+ \frac{2}{k^2} e^{-ik\mu\tau} \{h'' + 6\eta'' - ik\mu(h' + 6\eta')\}_{\tau=0}. \end{aligned} \quad (\text{B.5})$$

The last bracket contains the initial value of $(h' + 6\eta')$ and $(h'' + 6\eta'')$. The former vanishes for all types of initial conditions excepted NIV; the latter is non-zero for AD, NID and NIV. All initial condition terms in the first and last lines can be grouped and represented by two coefficients α and β :

$$\alpha(\mathbf{k}) - i\mu\beta(\mathbf{k}) \equiv F^0(\mathbf{k}, \mu) + \frac{2}{k^2} \{h'' + 6\eta'' - ik\mu(h' + 6\eta')\}_{\tau=0} \quad (\text{B.6})$$

$$= \left\{ \delta_{\text{ur}} + \frac{2}{k^2} (h'' + 6\eta'') \right\}_{\tau=0} - \frac{4i\mu}{k} \left\{ \theta_{\text{ur}} + \frac{1}{2}(h' + 6\eta') \right\}_{\tau=0}, \quad (\text{B.7})$$

with e.g. $(\alpha, \beta) = (20/(15 + 4R_{\text{ur}}), 0)$ for adiabatic initial conditions (as can be checked from [20]), R_{ur} being the fractional contribution of ultra-relativistic species to the background density. It is now easy to expand the full solution in Legendre coefficients, using the definition $F(\mu) = \sum_l (-i)^l (2l + 1) F_l P_l(\mu)$ and the fact that plane waves can be expanded in terms of spherical Bessel functions:

$$\begin{aligned} F_l(\mathbf{k}, \tau) &= \alpha(\mathbf{k}) j_l(k\tau) + \beta(\mathbf{k}) j_l'(k\tau) \\ &+ \frac{2}{k^2} \int_0^\tau j_l(k(\tau - \tilde{\tau})) \{2k^2\eta'(\tilde{\tau}) + h'''(\tilde{\tau}) + 6\eta'''(\tilde{\tau})\} d\tilde{\tau} \\ &- \frac{2}{k^2} \left\{ (h'' + 6\eta'')\delta_{l0} + \frac{k}{3}(h' + 6\eta')\delta_{l1} \right\}. \end{aligned} \quad (\text{B.8})$$

The terms in the first line show how initial conditions propagate to later times, by just free-streaming. The other terms show how perturbations adjust themselves to the power injected in the system at any time by metric perturbations.

B.2 Sub-Hubble approximation

Well inside the Hubble radius, the above formal solution can be approximated by a simpler expression. The results of this subsection are never used in our UFA scheme or in CLASS, but we present them for completeness, and also because the approximation performed in the next subsection will follow the same logic.

The second line of the solution contains a convolution between a function which varies smoothly over a Hubble time (at least for $k\tau \gg 1$), and a Bessel function $j_l(x)$ with $x \equiv k(\tau - \tilde{\tau})$ which oscillates over $\tau_k = 1/k$. Bessel functions $j_l(x)$ peak near $x_{\text{peak}} = l + 1/2$ (in fact this statement is accurate only for very large l ; for instance, $j_1(x)$ peaks near $x_{\text{peak}} = 2$ and $j_2(x)$ near $x_{\text{peak}} = 3.5$). The integral on $\tilde{\tau} \in [0, \tau]$ corresponds to $x \in [0, k\tau]$. The goal of this subsection is to find an approximation for this convolution for small l values.

As long as $k\tau \leq 1$, it is difficult to find a low- l approximation for the convolution; the result is a function oscillating over a characteristic time $\tau_k = 1/k$. In this regime, the integral brings an extra oscillatory contribution to the term $\alpha(\mathbf{k})j_l(k\tau) + \beta(\mathbf{k})j'_l(k\tau)$; this explains while around Hubble crossing, the numerical solution for $F_l(\mathbf{k}, \tau)$ exhibits irregular oscillatory patterns, with very different peak amplitude between two consecutive periods.

However, when $k\tau \gg 1$, the integral runs over a large range $x \in [0, k\tau]$. For low l , this means that the convolution picks up significant contributions only near $x = x_{\text{peak}} \ll k\tau$. Near this value, the slowly-varying argument can be approximated as a constant, to be evaluated around $\tilde{\tau} = (k\tau - x_{\text{peak}})/k$, i.e. in very good approximation near τ , since $k\tau \gg x_{\text{peak}}$. So, we can write:

$$\begin{aligned} & \int_0^\tau j_l(k(\tau - \tilde{\tau})) \{2k^2\eta'(\tilde{\tau}) + h'''(\tilde{\tau}) + 6\eta'''(\tilde{\tau})\} d\tilde{\tau} \\ \longrightarrow & \{2k^2\eta'(\tau) + h'''(\tau) + 6\eta'''(\tau)\} \int_0^\tau j_l(k(\tau - \tilde{\tau})) d\tilde{\tau}. \end{aligned} \quad (\text{B.9})$$

Finally, still in this limit $k\tau \gg 1$, the last integral can be approximated by

$$\frac{1}{k} \int_0^\infty j_l(x) dx = \frac{\sqrt{\pi}\Gamma(l/2 + 1/2)}{2k\Gamma(l/2 + 1)} \quad (\text{B.10})$$

A more accurate approximation scheme would lead to extra contributions involving time derivatives of the quantity between brackets in (B.9): for instance, the next order term would be of the type

$$\{2k^2\eta''(\tau) + h^{(4)}(\tau) + 6\eta^{(4)}(\tau)\} \frac{\gamma}{k^2} \quad (\text{B.11})$$

with γ being a coefficient of order one. However, since inside the Hubble scale metric perturbations vary over a Hubble time $\tau \gg \tau_k$, we can keep only the leading source terms with the highest power of k :

$$\delta_{\text{ur}} = F_0 = \alpha j_l(k\tau) + \beta j'_l(k\tau) + \frac{4}{k} \sqrt{\frac{\pi}{2}} \eta', \quad (\text{B.12})$$

$$\frac{4}{3k} \theta_{\text{ur}} = F_1 = \alpha j_l(k\tau) + \beta j'_l(k\tau) - \frac{2}{3k} h', \quad (\text{B.13})$$

$$2\sigma_{\text{ur}} = F_2 = \alpha j_l(k\tau) + \beta j'_l(k\tau) + \frac{\pi}{k} \eta'. \quad (\text{B.14})$$

In the code, we do not use directly these asymptotic approximations, and try instead to close the system of differential equations with a truncation at order two, like for a viscous fluid.

B.3 Exact truncation formula

First, in order to manipulate more compact equations, we define:

$$G_l(\mathbf{k}, \tau) \equiv F_l(\mathbf{k}, \tau) - \frac{2}{k^2} \left\{ (h'' + 6\eta'')\delta_{l0} + \frac{k}{3}(h' + 6\eta')\delta_{l1} \right\}. \quad (\text{B.15})$$

Since $j'_l(x) = j_{l-1}(x) - \frac{l+1}{x}j_l(x)$ and G_l contains a term $j_l(k\tau)$, we suspect that it obeys approximately to a similar relation. We use equation (B.8) for computing exactly the difference $G'_l - kG_{l-1} + \frac{l+1}{\tau}G_l$, which would vanish if only the term $\alpha j_l(k\tau)$ was contributing. We use the identities

$$j''_l(x) = j'_{l-1}(x) - \frac{l+1}{x}j'_l(x) + \frac{l+1}{x^2}j_l(x) \quad (\text{B.16})$$

and $j_l(0) = \delta_{l0}$. We find

$$\begin{aligned} G'_l - kG_{l-1} + \frac{l+1}{\tau}G_l &= \beta \frac{l+1}{k\tau^2} j_l(k\tau) + \frac{2}{k^2} \{2k^2\eta'(\tau) + h'''(\tau) + 6\eta'''(\tau)\} \delta_{l0} \\ &\quad + \frac{2}{k^2} \int_0^\tau K(\tau, \tilde{\tau}) \{2k^2\eta'(\tilde{\tau}) + h'''(\tilde{\tau}) + 6\eta'''(\tilde{\tau})\} d\tilde{\tau} \end{aligned} \quad (\text{B.17})$$

where we defined

$$K(\tau, \tilde{\tau}) \equiv k j'_l(k(\tau - \tilde{\tau})) - k j_{l-1}(k(\tau - \tilde{\tau})) + \frac{l+1}{\tau} j_l(k(\tau - \tilde{\tau})) \quad (\text{B.18})$$

This expression simplifies to

$$K(\tau, \tilde{\tau}) = -\frac{l+1}{\tau} \left(\frac{\tilde{\tau}}{\tau - \tilde{\tau}} \right) j_l(k(\tau - \tilde{\tau})) \quad (\text{B.19})$$

This truncation scheme is not practical in general. However, the goal of the UFA is to find a way to close the system at low l not at all times, but only deep inside the Hubble radius.

B.4 Sub-Hubble truncation formula

In the limit $k\tau \gg 1$, equation (B.17) can be simplified for two reasons. First, the Bessel function varies over a time scale $\tau_k = 1/k \ll \tau$, so

$$j'_l(k\tau) \gg \frac{j_l(k\tau)}{k\tau} \quad (\text{B.20})$$

Hence, in this limit, the last term in the identity (B.16) can be omitted, which implies that the first term (proportional to β) in eq. (B.17) is always negligible, even in the presence of isocurvature modes. Second, we can devise an approximation for the convolution, by noticing once more that it involves metric perturbations which vary smoothly over a Hubble time (at least for $k\tau \gg 1$), and the quantity $\frac{k\tau-x}{x}j_l(x)$ with $x \equiv k(\tau - \tilde{\tau})$ which oscillates over a period of order $\tau_k = 1/k$.

When $k\tau \gg 1$, the integral runs over a large range $x \in [0, k\tau]$. For low l , this means that the convolution picks up significant contributions only near $x = x_{\text{peak}} \ll k\tau$, while in this range $\frac{k\tau-x}{x}j_l(x) \simeq \frac{k\tau}{x}j_l(x)$. Near $\tilde{\tau} = (k\tau - x_{\text{peak}})/k \simeq \tau$, the slowly-varying metric perturbations can be treated as a constant term. So, we can write:

$$\begin{aligned} &\int_0^\tau \left(\frac{\tilde{\tau}}{\tau - \tilde{\tau}} \right) j_l(k(\tau - \tilde{\tau})) \{2k^2\eta'(\tilde{\tau}) + h'''(\tilde{\tau}) + 6\eta'''(\tilde{\tau})\} d\tilde{\tau} \\ &\longrightarrow \{2k^2\eta'(\tau) + h'''(\tau) + 6\eta'''(\tau)\} \int_0^\tau \left(\frac{\tilde{\tau}}{\tau - \tilde{\tau}} \right) j_l(k(\tau - \tilde{\tau})) d\tilde{\tau} \end{aligned} \quad (\text{B.21})$$

with

$$\int_0^\tau \left(\frac{\tilde{\tau}}{\tau - \tilde{\tau}} \right) j_l(k(\tau - \tilde{\tau})) d\tilde{\tau} \simeq \tau \int_0^\infty x^{-1} j_l(x) dx. \quad (\text{B.22})$$

Let us estimate the error made in these approximations. Since we are inside the Hubble radius with smoothly varying metric perturbations, and since the integral in eq. (B.21) is of order τ , the leading term in eq. (B.21) is of order $(k^2\tau\eta')$. The other terms of order $(\tau\eta''')$ and $(\tau h''')$ can be neglected. If instead of considering $\eta'(\tilde{\tau})$ as a constant we perform a Taylor expansion of this function around τ , we find that the next order contribution to eq. (B.21) is of order $(k\tau\eta'')$. Finally, an explicit calculation shows that the approximation performed in (B.22) amounts in neglecting terms of order k^{-1} with respect to terms of order τ . In summary, we obtain the following approximate truncation equation:

$$G'_l - kG_{l-1} + \frac{l+1}{\tau}G_l = \frac{2}{k^2} \{2k^2\eta'(\tau) + h'''(\tau) + 6\eta'''(\tau)\} \delta_{l0} - 4(l+1)\eta' \int_0^\infty \frac{j_l(x)}{x} dx + \mathcal{O}\left(\frac{\eta''}{k}, \frac{\eta'}{k\tau}\right). \quad (\text{B.23})$$

For $l = 2$, the integral is equal to $-1/3$, since $(j_1(x)/x)' = j_2(x)/x$ and $\lim_{x \rightarrow 0}[j_1(x)/x] = \frac{1}{3}$. So,

$$G'_2 - kG_1 + \frac{3}{\tau}G_2 = 4\eta' + \mathcal{O}\left(\frac{\eta''}{k}, \frac{\eta'}{k\tau}\right). \quad (\text{B.24})$$

Replacing the G_l 's by the appropriate momenta, we find:

$$\sigma'_{\text{ur}} = -\frac{3}{\tau}\sigma_{\text{ur}} + \frac{2}{3}\theta_{\text{ur}} - \frac{1}{3}(h' + 6\eta') + 2\eta' + \mathcal{O}\left(\frac{\eta''}{k}, \frac{\eta'}{k\tau}\right). \quad (\text{B.25})$$

The last term $2\eta'$ comes from our approximation for the convolution. Without a full treatment like the one presented here, one would miss this term and obtain the truncation formula called `ufa_mb` in the code, based on simply assuming $G_l \propto j_l(k\tau)$. However, with this extra contribution, the two terms in η' cancel each other, and at leading order we end up with

$$\sigma'_{\text{ur}} = -\frac{3}{\tau}\sigma_{\text{ur}} + \frac{2}{3}\theta_{\text{ur}} - \frac{1}{3}h', \quad (\text{B.26})$$

which is precisely what we call the `ufa_class` truncation scheme in `CLASS`.

References

- [1] E. Bertschinger, *COSMICS: cosmological initial conditions and microwave anisotropy codes*, [astro-ph/9506070](#) [[SPIRES](#)].
- [2] U. Seljak and M. Zaldarriaga, *A line of sight approach to cosmic microwave background anisotropies*, *Astrophys. J.* **469** (1996) 437 [[astro-ph/9603033](#)] [[SPIRES](#)].
- [3] A. Lewis, A. Challinor and A. Lasenby, *Efficient computation of CMB anisotropies in closed FRW models*, *Astrophys. J.* **538** (2000) 473 [[astro-ph/9911177](#)] [[SPIRES](#)].
- [4] M. Doran, *CMBEASY: an object oriented code for the cosmic microwave background*, *JCAP* **10** (2005) 011 [[astro-ph/0302138](#)] [[SPIRES](#)].
- [5] J. Hamann, A. Balbi, J. Lesgourgues and C. Quercellini, *Optimising Boltzmann codes for the Planck era*, *JCAP* **04** (2009) 011 [[arXiv:0903.0382](#)] [[SPIRES](#)].

- [6] J. Lesgourgues, *The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview*, [arXiv:1104.2932](#) [SPIRES].
- [7] J. Lesgourgues and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS) IV: Efficient implementation of non-cold relics*, [arXiv:1104.2935](#) [SPIRES].
- [8] J. Lesgourgues, *The Cosmic Linear Anisotropy Solving System (CLASS) III: Comparison with CAMB for LambdaCDM*, [arXiv:1104.2934](#) [SPIRES].
- [9] P.J.E. Peebles and J.T. Yu, *Primeval adiabatic perturbation in an expanding universe*, *Astrophys. J.* **162** (1970) 815 [SPIRES].
- [10] C.-P. Ma and E. Bertschinger, *Cosmological perturbation theory in the synchronous and conformal Newtonian gauges*, *Astrophys. J.* **455** (1995) 7 [[astro-ph/9506072](#)] [SPIRES].
- [11] M. Doran, *Speeding up cosmological Boltzmann codes*, *JCAP* **06** (2005) 011 [[astro-ph/0503277](#)] [SPIRES].
- [12] F.-Y. Cyr-Racine and K. Sigurdson, *Photons and baryons before atoms: improving the tight-coupling approximation*, *Phys. Rev. D* **83** (2011) 103521 [[arXiv:1012.0569](#)] [SPIRES].
- [13] W. Hu, *Structure formation with generalized dark matter*, *Astrophys. J.* **506** (1998) 485 [[astro-ph/9801234](#)] [SPIRES].
- [14] C. Pitrou, *The tight-coupling approximation for baryon acoustic oscillations*, *Phys. Lett. B* **698** (2011) 1 [[arXiv:1012.0546](#)] [SPIRES].
- [15] S. Naoz and R. Barkana, *Growth of linear perturbations before the era of the first galaxies*, *Mon. Not. Roy. Astron. Soc.* **362** (2005) 1047 [[astro-ph/0503196](#)] [SPIRES].
- [16] A. Lewis, *Linear effects of perturbed recombination*, *Phys. Rev. D* **76** (2007) 063001 [[arXiv:0707.2727](#)] [SPIRES].
- [17] L.F. Shampine and M.W. Reichelt, *The MATLAB ODE suite*, *SIAM J. Sci. Comput.* **18** (1997) 1.
- [18] T.A. Davis, *Direct methods for sparse linear systems*, SIAM, Philadelphia U.S.A. (2006).
- [19] M. Shoji and E. Komatsu, *Massive neutrinos in cosmology: analytic solutions and fluid approximation*, *Phys. Rev. D* **81** (2010) 123516 [Erratum *ibid.* **D 82** (2010) 089901] [[arXiv:1003.0942](#)] [SPIRES].
- [20] M. Bucher, K. Moodley and N. Turok, *The general primordial cosmic perturbation*, *Phys. Rev. D* **62** (2000) 083508 [[astro-ph/9904231](#)] [SPIRES].