Introduction

Though the natural world is an interconnected whole, our models usually treat small pieces of that whole as if they were isolated from everything else. For instance, in an elementary physics class, we are used to conservation of energy and momentum for *isolated* systems. Isolation can be more than just physical: we are used to the notion of *separation* between time and/or spatial scales. The flow around a body inherits a characteristic length from the body itself. At a typically much smaller characteristic distance from the body's surface, the flow is modified strongly by the action of viscous stresses. This is of course the archetypal example of a *boundary layer*. In physical systems exhibiting oscillations, the properties of the oscillation sometimes change in a characteristic time much longer than the local period, and this slow cumulative change is called *modulation* of the oscillations.

The mathematical models of phenomena with widely separated scales are often *singularly perturbed*, meaning that the solution of the equations doesn't converge uniformly as the ratio of scales becomes large or small. This is the conceptual frame of singular perturbation theory. As a practice or art, singular perturbation theory is a body of analysis that exploits the separation of scales in phenomena: First, describe small and large scale happenings as if isolated or separate from each other. Then join them so they talk to each other, and larger meanings emerge. This characterization of singular perturbation theory begs a question: What happens if there is a hierarchy of many scales with no clear large and small separation? Like turbulence. Maybe this larger sphere of problems is the future of the subject. In any case, it is beyond the pay grade of this book. So here is what this book *is* going to do:

First and generally, like the preceding text by the author, *Training Manual on Transport and Fluids*, there is a main text of basic material and a subtext of worked problems that go deeper and present engaging examples. Think of the main text as the trunk of the tree and the problems as the branches (with many bifurcations, as you will see).

Chapter 1 is a traditional introduction based on simple, preferably exactly solvable, examples of singular perturbation. We gain first impressions of scaling, dominant balances, distinguished limits, boundary layers, matching, and modulated oscillations.

All perturbation analysis is approximation, and Chapter 2 spells out the specific sense of *asymptotic* approximation in which "the error is much smaller than the smallest term we keep," as the perturbation parameter goes to zero. There is a brief overview of a very traditional subject: the asymptotic expansion of integrals. The long-time analysis of the Fourier integrals representing wavefields is a jumping-off point for a mini-course on WKB at the end of the chapter.

Chapter 3 is a traditional presentation of matched asymptotic expansions for ODE boundary value problems with localized small scale structure, such as boundary layers, internal layers, derivative layers, etc. The discussion of higher-order matching and how the overlap domain shrinks as the order of matching increases is based on the intention to make it the simplest possible, but no simpler.

Chapter 4 on moving internal layers introduces Chapman–Enskog asymptotics: There are dynamical systems in which a relatively small set of state variables dominates the solution. If you know the evolutions of these dominant variables, you know the evolution of the whole system. The flow vector field of the dominant variables is to be determined as an asymptotic expansion. Also, we construct asymptotic expansions of the other non-dominant variables, taking the dominant variables as given. The solvability conditions encountered in constructing the latter expansions dictate, order by order, the flow vector field of the dominant variables. In the first Chapter 4 example, the dominant variable is the centerline curve of the internal layer of a director field in two dimensions, and the asymptotic construction of the internal layer solution about this curve dictates, order by order, the dynamics of the centerline curve. At leading order, we get the familiar motion by curvature. The remaining content of Chapter 4 on projected Lagrangians is easiest to discuss in the context of Chapter 6.

Chapter 5 is a (nearly) traditional presentation of the Prandtl boundary layer theory for the Navier-Stokes equations and of the solutions of the boundary layer equations that can be constructed because of scaling symmetry. The most original part of this chapter is the last problem (Problem 5.5) on a spiral diffusion layer in a vortex flow.

Chapter 6 is the first attempt at analysis of modulated oscillations. It starts with an expansive repertoire of elementary examples (Problems 6.1–6.6). Even though their methodology is extremely simple (elementary exact solutions, WKB, perturbed ODE eigenvalue problem), we can present engaging examples, such as passage through resonance. The mathematical technique in the body of this chapter is the *method of two scales.* The literature often refers to the *multiple* scale method, because we might want to consider more than two characteristic times explicitly. It is this author's belief that the method of two scales is an introductory method, to be eventually superseded by averaging and its big brothers, which we introduce in Chapter 7. Once you are in the realm of these methods, the need for more than two characteristic times is moot. In Chapter 6, we'll see how the two scale analysis of nonlinear oscillations leads to the insight that the *action* in the sense of classical mechanics *is* the proper variable of modulation theory. In this sense, the two scale analysis is a precursor to the methods related to averaging in Chapter 7. These methods start with action as a state variable right away. The discovery of action by the method of two scales derives from Whitham's analysis of nonlinear waves, so the Whitham modulation theory of waves is a core subject of Chapter 6. Finally, there is Whitham's packaging of modulation theory for nonlinear variational equations by means of the *averaged Lagrangian.* Our main use of the averaged Lagrangian happens in two places: here in Chapter 6, we apply it to the homogenization theory of the effective diffusion tensor in a periodic medium. In the last problem (Problem 6.15), it is shown that the Lagrangian flavor of homogenization used in this example is equivalent to the traditional direct analysis. The projected Lagrangian in Chapter 4 is essentially averaging the original full Lagrangian over the internal layer. It is just like the Whitham analysis, except that the wave has only one crest.

Chapter 7 is about modulation theory of a perturbed Hamiltonian dynamics with one degree of freedom, based on perturbation of its action-angle variables. This is a special case of the more general class of problems treated by the *averaging method.* Our focus is a bit narrow in the interest of staying clear and simple in a textbook. The essential idea: When a perturbation is applied to the original Hamiltonian dynamics, the action variable tends to undergo large, slow drifts, with a small-amplitude, rapid oscillation superpositioned on top of it. The idea in this chapter is to perform a near-identity transformation of the original action-angle variables, so the new action has no rapid, oscillatory component. In a process very reminiscent of the Chapman–Enskog method of Chapter 4, the governing ODE of the new action and the small oscillatory correctors in the near-identity transformation have intertwined asymptotic expansions. If we restrict the analysis to leading order, the resulting modulation theory is completely equivalent to the well-known averaging method. As we proceed to higher order, we make ever more refined corrections to the dynamics of the new action, and it is hoped that the characteristic time of validity of the asymptotics is increased. This is why this book has not pursued multiple scale theory with more than two characteristic times. The chapter closes with an analysis of dissipative perturbations of the Kepler problem. Those of you who have donated a coin to the gravity well exhibit in a planetarium will appreciate the result of this analysis: you'll know why the orbit of your dime is almost circular just before it spirals into the "black hole" of the donation box.

Chapter 8 introduces into the perturbed Hamiltonian dynamics a feature that is expressly avoided in Chapter 7, and that is explicit periodic time dependence of the perturbation. Why is that a big deal? If the frequency of the unperturbed Hamiltonian orbit is sufficiently close to a rational multiple of the perturbation frequency, formal asymptotics as in Chapter 7 predicts *resonance*. A sure sign of resonance is deviations from the unperturbed orbit that *don't* scale in direct proportion to the perturbation. They are much larger, and their characteristic time is much longer than the perturbation frequency. The perplexing issue is that the rational numbers are dense, so the ability of simple asymptotics to isolate one resonance at a time seems dubious. Chapter 8 carries out the obvious program: Just *do* the simple resonance asymptotics anyway, and compare with direct numerical solutions of the full ODE. In the elementary cases examined, the asymptotics displays clear robustness within the formal order of approximation and over characteristic times for which the asymptotics is valid. At the end of Chapter 8 we "look through a glass darkly" by means of a simple formal estimate: The resonance associated with a given rational frequency ratio is felt in a narrow band of the phase plane about a given unperturbed orbit. As the strength of the perturbation decreases, so does the bandwidth. If the resonance associated with the rational number M/N wants to be in the bandwidth of another resonance, say M^*/N^*, M and N generally go to infinity as the difference from M^*/N^* goes to zero. If M and N are large, the formal perturbation theory shows that the bandwidth and strength of the M/N resonance goes to zero. In summary, when M/N is close enough so its resonance is in the bandwidth of the M^*/N^* resonance, the latter resonance can't resolve it. The great and perplexing questions about resonance at the level of rigorous analysis remain, and no claim is made here in relation to them.