Published in final edited form as:

J Phys Chem B. 2009 February 12; 113(6): 1640–1644. doi:10.1021/jp806431b.

The Deterministic Limit of Stochastic Chemical Kinetics

Daniel T. Gillespie

Dan T. Gillespie Consulting, 30504 Cordoba Pl., Castaic, California 91384

Abstract

An analysis is presented of the approximating assumptions that underlie a recently proposed derivation of the traditional deterministic reaction rate equation from a discrete-stochastic formulation of chemical kinetics. It is shown that if the system is close enough to the thermodynamic limit, in which the molecular populations and the containing volume all approach infinity in such a way that the molecular concentrations remain finite, then the required approximating assumptions will be justified for practically all spatially homogeneous systems that one is likely to encounter.

I. Introduction

The transitioning of spatially homogeneous chemical kinetics from its discrete-stochastic foundations in molecular physics to the traditional continuous-deterministic formalism of ordinary differential equations can be viewed as occurring in stages, as certain conditions become satisfied.¹ As sketched in Fig. 1, the starting point for this scenario is the hypothesis that each reaction channel is described by a propensity function, or stochastic reaction rate. This is a function of the current molecular populations whose product with infinitesimal time dt gives the probability that the reaction will occur somewhere inside the system in the next dt . This definition leads exactly, via the laws of probability, to both the chemical master equation (CME) and the stochastic simulation algorithm (SSA).¹ But, as will be described in more detail below, under certain conditions this definition also leads, albeit approximately, to the Poisson tau-leaping formula. That in turn approximates, under certain additional conditions, to the chemical Langevin equation (CLE). And the CLE finally reduces, in the thermodynamic limit, to the classical deterministic reaction rate equation (RRE). The thermodynamic limit that enables the last step is formally defined as the limit in which the molecular populations and the containing volume approach infinity together, and in such a way that the molecular concentrations (the ratios of the molecular populations to the system volume) remain finite.

What is missing from this narrative is an explanation of how the logical connections among the assumptions that enable the three approximating steps actually conspire to make the argument logically inevitable, at least for the kinds of chemical reactions that we may expect to encounter in practice. This is far from obvious, because the conditions that enable the first two steps appear to be potentially incompatible. In this paper we will show why that potential incompatibility does not, in fact, pose a problem for the scenario sketched in Fig. 1.

We should note that there have been a number of earlier efforts aimed at understanding how and under what circumstances a continuous, deterministic description in terms of ordinary differential equations can be achieved for a dynamical system that is intrinsically discrete and stochastic. Perhaps the most important of these is the seminal work of Kurtz², and the alternative approach of van Kampen³. While not questioning the correctness of those earlier analyses, it seems fair to say that the arguments they employ are mathematically dense, and

for many chemists and physicists they do not convey a clear sense of what is going on physically. The derivation described in this paper attempts to keep the mathematics as simple as possible, and to bring underlying physical ideas to the fore.

II. The First Two Steps

We begin with a brief review of the first two approximating steps in the progression outlined in Fig. 1. (A fuller discussion can be found in Ref. 1 and other references therein.) The propensity function a_j of reaction channel R_j is defined so that, if the system at time t is in the molecular population state $\mathbf{x}$, then $a_j(\mathbf{x})dt$ gives the probability that reaction R_j will fire in the next infinitesimal time interval $[t, t + dt)$. The key assumption is that, for any chemical system that is kept spatially homogeneous or “well-stirred”, this definition accurately describes how chemical reaction events actually occur. If R_j does fire when the system is in state $\mathbf{x}$, the system immediately jumps to state $\mathbf{x} + \mathbf{v}_j$, where $\mathbf{v}_j$ is called the state-change vector of R_j . Together, the quantities $\{a_j\}$ and $\{\mathbf{v}_j\}$ completely determine the dynamical behavior of the system.

Given that the system at time t is in some state $\mathbf{X}(t) = \mathbf{x}$, let the integer random variable $n_j(\mathbf{x}, \Delta t)$ denote the number of times reaction R_j will fire in the finite time interval $[t, t + \Delta t)$. Then the net change in state in $[t, t + \Delta t)$ will be $\sum_j n_j(\mathbf{x}, \Delta t) \mathbf{v}_j$. Now suppose that Δt is chosen small enough that none of the propensity functions changes its value appreciably during $[t, t + \Delta t)$. Then it follows from the definition of the propensity function and the definition of the Poisson random variable that $n_j(\mathbf{x}, \Delta t)$ can be approximated by $\mathcal{P}_j(a_j(\mathbf{x})\Delta t)$, the Poisson random variable with mean and variance $a_j(\mathbf{x})\Delta t$. The system's state at time $t + \Delta t$ can then be *approximately* computed as

$$\mathbf{X}(t + \Delta t) \doteq \mathbf{x} + \sum_j \mathcal{P}_j(a_j(\mathbf{x})\Delta t) \mathbf{v}_j. \quad (1)$$

This equation is called the Poisson tau-leaping formula. The condition justifying it, namely that none of the propensity functions changes “appreciably” during $[t, t + \Delta t)$, is called the Leap Condition. The Leap Condition is the basis for the first approximating step in Fig. 1. It essentially requires Δt to be small enough that the change $\Delta a_j(\mathbf{x})$ in each propensity function during $[t, t + \Delta t)$ satisfies⁴

$$\left| \frac{\Delta a_j(\mathbf{x})}{a_j(\mathbf{x})} \right| \ll 1. \quad (2)$$

The second approximating step in Fig. 1 makes the assumption that Δt , although *small* enough to satisfy condition (2), can nevertheless be chosen *large* enough that

$$a_j(\mathbf{x})\Delta t \gg 1. \quad (3)$$

It is well known that the Poisson random variable with mean and variance m can, when $m \gg 1$, be approximated by the *normal* random variable with the same mean and variance. Therefore, when (3) holds, each Poisson random variable in Eq. (1) can be approximated as

$$\mathcal{P}_j(a_j(\mathbf{x})\Delta t) \doteq \mathcal{N}_j(a_j(\mathbf{x})\Delta t, a_j(\mathbf{x})\Delta t) \equiv a_j(\mathbf{x})\Delta t + \sqrt{a_j(\mathbf{x})\Delta t} \mathcal{N}_j(0,1),$$

where

$\mathcal{N}$

$_j(0,1)$ is the normal random variable with mean 0 and variance 1, and the last step is an elementary result in random variable theory. This approximation transforms the tau-leaping formula (1) into the chemical Langevin equation (CLE) ¹:

$$\mathbf{X}(t+\Delta t) \doteq \mathbf{x} + \sum_j \nu_j a_j(\mathbf{x})\Delta t + \sum_j \nu_j \sqrt{a_j(\mathbf{x})} \mathcal{N}_j(0,1) \sqrt{\Delta t}. \quad (4)$$

But here a critical question arises: How do we know that there will exist a time interval Δt that is at once *small* enough to satisfy condition (2) yet also *large* enough to satisfy condition (3)? Or from a more practical point of view, under what conditions can we be assured that the system will be blessed with such an accommodating Δt ?

III. Real Chemical Reactions

In Sec. IV we will prove that, for the kinds of chemical reactions normally encountered in practice, it will always be possible to find a Δt that satisfies *both* conditions (2) and (3), *provided the system is sufficiently close to the thermodynamic limit*. As mentioned earlier, the thermodynamic limit is defined as the limit in which the reactant populations x_i and the system volume Ω all become infinitely large, but in such a way that the reactant concentrations $x_i/\Omega \equiv z_i$ stay fixed. It is important to understand that the thermodynamic limit is not a limit that the system actually approaches as a consequence of its natural temporal evolution, nor even as a result of the experimenter's direct intervention; rather, the thermodynamic limit is a kind of *gedanken limit*, an idealized state that turns out to be useful because it often provides a convenient approximation to macroscopic systems.

Our derivation in Sec. IV will depend critically on some special properties of “real world” chemical reactions that have not yet been mentioned. To describe those properties and explain briefly their origins, we shall henceforth let a denote a generic propensity function, x the molecular population of a generic reactant species, and $z = x/\Omega$ a generic species concentration.

The first special property of real-world chemical reactions that we will make use of is this: The propensity functions for all chemical reactions that are encountered in practice behave in the thermodynamic limit as

$$a \sim \gamma x, \quad (5)$$

where γ is a constant with respect to the thermodynamic limit, but not necessarily a constant with respect to time. (We use the symbol $\sim$ to denote equality in the thermodynamic limit.) Property (5) says that, in the thermodynamic limit, every propensity function a *diverges linearly* with the “size” of the system; or equivalently, a/Ω *remains constant* in the thermodynamic limit.

The reason why (5) is true depends on whether the reaction in question is unimolecular (first order), bimolecular (second order), or source-like (zeroth order). In a *unimolecular* reaction, a molecule of a particular species spontaneously undergoes some prescribed chemical change.

For any such reaction encountered in practice, it is an empirical fact that, on time scales of chemical interest, there will exist some temporal constant c such that $c dt$ gives the probability that a particular molecule will undergo that reaction in the next infinitesimal time dt . This constant c depends on structural (often quantum mechanical) properties of the molecule, and possibly also on the system temperature; however, it will not depend on the system volume Ω nor the number of molecules in the system. If there are x molecules of this species in the system, then by the addition law of probability it follows that the probability that some one of them will so react in the next dt is $x \times c dt = (cx)dt$. The coefficient of dt in this last probability is, by definition, the reaction's propensity function, so we conclude that a unimolecular propensity function is exactly of the form (5) with $\gamma = c$.

In a *bimolecular* reaction, two molecules of the same or different species collide and react to produce some specific chemical change. The probability that a particular pair of reactant molecules will do that in the next dt can be computed as the product of {the probability that the two molecules will collide in the next dt } times {the probability that a colliding pair of molecules will react}. The first probability can in principle be computed from kinetic theory, treating the molecules as moving about in thermal equilibrium in either a dilute gas or a suspending diffusional medium. This collision probability typically takes the form¹ $(k'/\Omega)dt$, where k' depends on the collision cross section of the two molecules, and also on either their mean relative speed or their mutual diffusion coefficient; however, k' will not depend on the molecular populations or Ω . The inverse- Ω dependence of the collision probability arises roughly because in a well-stirred system, two molecules will be less likely to collide with each other when their confining volume is made larger. The probability of a reaction given a collision is often quite difficult to calculate from first principles, but for our purposes it suffices to note that this probability will be independent of the molecular populations and Ω . The product of the {collision probability in the next dt } and {the reaction probability given a collision} thus gives for the probability that a particular pair of molecules will react in the next dt an expression of the form $(k/\Omega)dt$, with k independent of both the molecular populations and Ω . Summing this reaction probability over all possible reactant pairs, which number $x_1 x_2$ if the reacting species 1 and 2 are different, or $\frac{1}{2}x(x-1)$ if they are the same, then gives the final probability that the reaction will happen somewhere in the system in the next dt . In the thermodynamic limit of large x and Ω , this final probability thus takes the form

$$\left(\frac{k}{\Omega}dt\right)\alpha x^2 = \left(\alpha k \frac{x}{\Omega}\right)x dt = ((\alpha k z)x)dt,$$

where α is either 1 or $\frac{1}{2}$ according to whether the reacting species are different or the same. The coefficient of dt on the right is, by definition, the propensity function of the reaction. Since α , k , and the concentration z are all constants with respect to the thermodynamic limit (although z is not a temporal constant), the bimolecular propensity function therefore has the property (5) with $\gamma = \alpha k z$.

Higher order chemical reactions, such as trimolecular reactions, actually consist of sequences of unimolecular and bimolecular reactions, and therefore they need not be considered separately. But the chemistry of open systems often involves *source-like* reactions that inject new molecules into the system. Such a reaction is typically modeled by the dynamical rule that, for parameter r , $r dt$ gives the probability that, in the next dt , a new molecule will be introduced at some uniformly random location inside Ω . But if this injection process is to behave sensibly in the thermodynamic limit, then r must have the form $r = r_0 \Omega$, where r_0 is independent of Ω and the molecular populations; because, if r grew any slower (or faster) than linearly with Ω , then the number of molecules introduced into the system per unit time *per unit*

volume would approach zero (or infinity) in the thermodynamic limit. So, the propensity function for a source-like reaction will typically have the form

$$r=r_0\Omega \sim \left(r_0\frac{\Omega}{x}\right)x=\left(\frac{r_0}{z}\right)x,$$

where x can be the population of any always-present reactant species, and z is the corresponding concentration. Since both r_0 and z are constants with respect to the thermodynamic limit, then this propensity function has the property (5) with $\gamma = r_0/z$.

The second feature of real-world chemical reactions that will be relevant to our considerations here concerns their state-change vectors: The absolute change in any species population x caused by a single reaction firing is practically always either 0, 1 or 2, and thus is never greater than 3 (the exact small bounding integer in any specific instance can be determined by simply inspecting all the v_j). It follows that if $n(x,\Delta t)$ is the number of times a given reaction channel will fire in the next Δt , then the absolute change $|\Delta x|$ in the molecular population of any reactant species caused by firings of that reaction channel in time Δt will be bounded by $3n(x,\Delta t)$. Therefore, if there are M reaction channels, then we have the (very conservative) $|\Delta x|$ bound

$$|\Delta x| < 3Mn(x,\Delta t). \quad (6)$$

Of course, we can easily imagine bizarre reactions for which properties (5) or (6) are *not* true. For such reactions the following analysis, and quite likely its conclusions, will not hold. But such reactions typically do not occur in real-world chemical systems.

IV. Derivation of the Classical Limit

Armed with properties (5) and (6), we shall now prove that *if the system is close enough to the thermodynamic limit, we can find a time interval Δt that satisfies both conditions (2) and (3).* As can be seen from Fig. 1, this would guarantee that the CLE (4) will hold whenever the system is, roughly speaking, “sufficiently large”.

We begin by taking the system close enough to the thermodynamic limit (i.e., we take both x and Ω large enough) that we can invoke property (5) to write the quantity on the left hand side of condition (2) as

$$\left|\frac{\Delta a}{a}\right| \sim \left|\frac{\Delta(\gamma x)}{\gamma x}\right| = \frac{|\Delta x|}{x} < \frac{3Mn(x,\Delta t)}{x}, \quad (7)$$

where the last step invokes property (6). Since the probability $adt \sim \gamma x dt$ that the reaction channel will fire in the next infinitesimal time dt goes to zero as $dt \rightarrow 0$, then we can make $n(x,\Delta t)$ arbitrarily small by taking Δt small enough. It then follows from (7) that, by taking Δt small enough, we can make $|\Delta a/a|$ small enough that we can invoke the earlier described Poisson approximation that leads to the tau-leaping formula (1), namely, $n(x,\Delta t) \approx \mathcal{P}(a\Delta t) \sim \mathcal{P}(\gamma x\Delta t)$. That approximation implies that

$$\langle n(x,\Delta t) \rangle \sim \gamma x\Delta t. \quad (8)$$

Since the Poisson random variable has a well defined (finite) standard deviation, there will exist a $B < \infty$ such that $n(x, \Delta t)$ can, at least for all practical purposes, be bounded by B times its mean:

$$n(x, \Delta t) < B \langle n(x, \Delta t) \rangle. \quad (9)$$

We will shortly show that in the thermodynamic limit, the value $B=2$ will easily suffice. Inserting (8) into (9), and the result into (7), we obtain

$$\left| \frac{\Delta a}{a} \right| < \frac{3M(B\gamma x \Delta t)}{x} = 3BM\gamma \Delta t. \quad (10)$$

The result (10) implies that we can assure condition (2), in the form

$$\left| \frac{\Delta a}{a} \right| < \varepsilon \quad (11)$$

for any given ε satisfying $0 < \varepsilon \ll 1$, by taking

$$\Delta t = \frac{\varepsilon}{3BM\gamma}. \quad (12)$$

Note that this value for Δt is *independent of* x . With Δt so chosen, it then follows that the quantity on the left hand side of condition (3) can be written

$$a\Delta t \sim \gamma x \Delta t = \gamma x \frac{\varepsilon}{3BM\gamma} = \frac{\varepsilon}{3BM} x. \quad (13)$$

This tells us that we can now satisfy condition (3) by taking

$$x \gg \frac{3BM}{\varepsilon}. \quad (14)$$

So we see that, for any given $\varepsilon > 0$, we can satisfy condition (2), in the form (11), and also condition (3), simply by choosing Δt and x according to (12) and (14). To complete our proof, we need only show that the bound B on $n(x, \Delta t)$ assumed in (9) is finite for arbitrarily large x . This we can do by noting that, since condition (2) implies that $n(x, \Delta t) \approx \mathcal{P}(a\Delta t)$, then the fact that the variance of the Poisson random variable is equal to its mean implies that

$$\langle n(x, \Delta t) \rangle \approx a\Delta t \text{ and } \text{sdev} \{n(x, \Delta t)\} \approx \sqrt{a\Delta t}. \quad (15)$$

But (13) tells us that, for x sufficiently large, $a\Delta t \gg 1$. So if, for example, $a\Delta t = 100$, and hence $\langle n(x, \Delta t) \rangle \approx 100$ and $\text{sdev} \{n(x, \Delta t)\} \approx 10$, then in order for $n(x, \Delta t)$ to exceed its mean by a factor

of $B=2$, $n(x, \Delta t)$ would have to experience a fluctuation of *ten* standard deviations. That is very unlikely. If $a\Delta t$ were increased from 100 to 400, the same logic shows that a fluctuation of *twenty* standard deviations in $n(x, \Delta t)$ would be needed in order for $n(x, \Delta t)$ to exceed twice its mean, and that can be regarded as “impossible” from a practical point of view. These considerations show that B in the above analysis can safely be set to the eminently finite value 2. This completes our proof that the CLE (4) will inevitably hold if the system is sufficiently close to the thermodynamic limit.

The forgoing argument can be briefly summarized as follows: When the reactant populations are sufficiently large in the sense of the thermodynamic limit, we can satisfy condition (2) by choosing for Δt a value that, although small, is independent of how large the reactant populations are. That done, we can then satisfy condition (3), without compromising condition (2), by taking the reactant populations sufficiently larger.

The final step from the CLE to the RRE in the full thermodynamic limit (see Fig. 1) is a simple consequence of property (5), namely, that in the thermodynamic limit the propensity functions *diverge linearly* with the system size. That implies that as we approach the thermodynamic limit, all terms in the CLE (4) grow linearly with x except for the last (stochastic) term, which grows more slowly as the square root of x ; thus, in the full thermodynamic limit, the stochastic term in (4) becomes vanishingly small compared to the other terms, and the CLE accordingly reduces to

$$\mathbf{X}(t+\Delta t) \doteq \mathbf{x} + \sum_j \nu_j a_j(\mathbf{x}) \Delta t. \quad (16a)$$

This is essentially the deterministic RRE. A more familiar form of the RRE can be obtained by first dividing (4) through by Ω , and then defining $\tilde{a}_j(\mathbf{z}) \equiv \Omega^{-1} a_j(\mathbf{x})$, where $\mathbf{z} \equiv \mathbf{x}/\Omega \equiv \mathbf{X}(t)/\Omega$. The stochastic term in the resulting equation then goes to zero in the thermodynamic limit, yielding the familiar “concentration” form of the RRE:

$$\mathbf{Z}(t+\Delta t) \doteq \mathbf{z} + \sum_j \nu_j \tilde{a}_j(\mathbf{z}) \Delta t. \quad (16b)$$

This completes our derivation of the RRE. But we should note that, when converting Eqs. (16) into ordinary differential equations by taking the limit of “vanishingly small” Δt , we must respect the condition that Δt cannot be made smaller than dictated by formula (12). In other words, we must be content to regard the *smallest* Δt that simultaneously satisfies conditions (2) and (3) as being “infinitesimally” small. This proviso effectively sets a lower limit on the timescale of the description of the system that is provided by the RRE.

V. Conclusions

Previous descriptions¹ of the logical progression to the RRE sketched in Fig. 1 were cautiously circumscribed, asserting that the progression would occur only if the system were such that a time interval $\tau = \Delta t$ could be found that simultaneously satisfies the two conditions that lead to the CLE. We have here strengthened that result by showing that, for any “real-world” chemical system, such a time interval will inevitably exist *if the system is sufficiently large in the sense of the thermodynamic limit*. Conveniently, that condition is also what ensures the final step from the CLE to the RRE.

The universal validity of the RRE in the thermodynamic limit might be challenged on the grounds that, in systems with two or more stable states (defined as relative maximums of the

$t \rightarrow \infty$ solution of the CME), spontaneous transitions between those states, which *inevitably* occur in the stochastic theory, are *not allowed* by the deterministic RRE. The answer to this challenge is that, as the thermodynamic limit is approached, the mean transition times between the stable states all increase very rapidly, and will eventually exceed the lifetime of the universe. In that circumstance, the claim of the RRE that transitions between stable states “never” occur will be justified in a practical sense. The drawback of the RRE in this regard is that it cannot itself predict how large the system must be before the possibility of a stable state transition can safely be ignored. For that we must, at least in principle, look to the CME or the SSA. And we must be careful not to overlook the possibility of stable state transitions by using the RRE to describe the behavior of very small systems that are *not* close to the thermodynamic limit, such as for instance some cellular systems.

Finally, we must emphasize that our analysis here has *not* provided a convenient *practical* answer to the important question of when it is okay to simulate a chemical system using the RRE instead of the SSA. For that, the condition “sufficiently close to the thermodynamic limit” would have to be rendered much more explicitly, and connected in an unambiguous way to easily computable error bounds on the population trajectories for any given situation. Developing methods for doing that would constitute a major advance beyond the present asymptotic result. Until that feat is accomplished, the question of whether it is okay to use the RRE instead of the SSA can only be answered *post facto*: If, for a given system, the molecular population trajectories generated by the SSA are found to be adequately approximated by the smooth, deterministic trajectories computed from the RRE, with all deviations in the former about the latter being inconsequentially small, then it is okay to use the computationally more efficient RRE; otherwise, it is not.

Acknowledgements

I thank the Linda Petzold Group at UCSB, and also the Journal reviewers, for some very beneficial suggestions. Financial support for this work was provided by the California Institute of Technology through Consulting Agreement 102-1080890 pursuant to Grant R01GM078992 from the National Institute of General Medical Sciences, and through Contract 82-1083250 pursuant to Grant R01EB007511 from the National Institute of Biomedical Imaging and Bioengineering. Support was also provided by the University of California at Santa Barbara under Consulting Agreement 054281A20 pursuant to funding from the National Institutes of Health. The content of this work is solely the responsibility of the author, and does not necessarily reflect the official views of any of the aforementioned institutions.

References and Notes

1. Gillespie D. Annu Rev Phys Chem 2007;58:35. [PubMed: 17037977]
2. Kurtz TG. J Chem Phys 1972;57:2976. Math Progr Stud 1976;5:67. Stoch Proc Appl 1978;6:223.
3. van Kampen, NG. Stochastic Processes in Physics and Chemistry. North-Holland: 1992.
4. Since no reactant population can change by an amount smaller than 1, the minimum possible non-zero change in $a_j(\mathbf{x})$ is its reaction-rate constant, c_j . Therefore, to accommodate circumstances in which $a_j(\mathbf{x})$ is near zero, in practical calculations the Leap Condition requirement $|\Delta a_j(\mathbf{x})| \ll a_j(\mathbf{x})$ is softened by never requiring $a_j(\mathbf{x})$ to change by an amount smaller than c_j . But this softening of the Leap Condition will not be required in the large system limit that we are concerned with in this paper.

Figure 1.

(From Ref. 1.) The theory of stochastic chemical kinetics is premised on the definition of the propensity function in the top box, a definition that ultimately must look to molecular physics for its justification. The two other solid outlined boxes in the figure denote *exact* consequences of that definition: the chemical master equation (CME) and the stochastic simulation algorithm (SSA). Dashed boxes denote *approximate* consequences: the tau-leaping formula, the chemical Langevin equation (CLE), the chemical Fokker-Planck equation (CFPE), and the reaction rate equation (RRE). The bracketed condition by each dashed inference arrow gives the condition under which that approximation is warranted. The condition justifying tau-leaping is called the Leap Condition, and the condition justifying the RRE is called the thermodynamic limit. This paper shows that, for real-world chemical systems, the *thermodynamic limit alone* guarantees satisfaction of the other two conditions, and thus that the top-to-bottom progression in this figure inevitably occurs in that limit.