

A rigorous derivation of the chemical master equation

Daniel T. Gillespie

Research Department, Naval Weapons Center, China Lake, CA 93555, USA

It is widely believed that the chemical master equation has no rigorous microphysical basis, and hence no *a priori* claim to validity. This view is challenged here through arguments purporting to show that the chemical master equation is *exact* for any gas-phase chemical system that is kept *well stirred* and *thermally equilibrated*.

1. Introduction

We often find it convenient to define the “state” of a well stirred chemically reacting system as “the current number of molecules of each component species”. But if we do that, then we must accept the fact that this state cannot evolve with time as purely *deterministic* process. The reason is simply that a specification of the molecular populations does *not* constitute a specification of the positions and momenta of all the molecules in the system, and hence does not suffice to uniquely determine the future behavior of the system.

Early efforts to mathematically accommodate the intrinsically stochastic nature of chemically reacting systems were reviewed in the 1967 article by McQuarrie [1]. That article introduced many scientists to the principle mathematical features of what is nowadays called the *chemical master equation*. McQuarrie did not claim in that article to give a rigorous microphysical derivation of the chemical master equation. He simply posited some “transition probabilities” in loose analogy to the conventional chemical reaction rates, and then invoked a “detailed balance” heuristic to obtain master equations for some specific chemical systems; his main concern was to develop solutions to those master equations. But if one cannot demonstrate a rigorous microphysical basis for the chemical master equation, then one cannot accord its solutions any *a priori* validity. In fact, the prevailing view today seems to be that the chemical master equation does *not* have *a priori* validity, and that its physical fidelity can be assessed only through *a posteriori* comparisons of its predictions with the results of experiments – either real experiments or molecular dynamics experiments.

The aim of this paper is to show that this view of the chemical master equation does not do it justice. Offered here is a purportedly rigorous derivation of the chemical master equation for a chemically reacting gas-phase system that is kept *well stirred* and *in thermal equilibrium*.

2. Rigor and probability

In general, a rigorous derivation consists of first laying down a few self-evident premises, and then making a series of self-evident inferences that leads to the desired conclusion. The problem with this otherwise sanguine procedure is that reasonable people can disagree on whether something is “self-evident”. In order to confront this problem openly, we shall begin by considering a purportedly rigorous derivation of the expected result of a very simple physical experiment. This exercise will allow us to introduce the principal rules of inference that will be used in our subsequent derivation of the chemical master equation, and it will also provide an indication of the kind of premises that we shall lay down for that derivation. The experiment that we propose to analyze concerns the tossing of a pair of dice. Specifically, we want to assess the probability that the fair tossing of a fair pair of cubic dice will yield a face-up sum of 4.

By the word “probability” here, we mean quite specifically the *fraction of trials* (in this case tosses) that yield a particular *event* (in this case a face-up sum of 4), in the limit of infinitely many repeated trials. This “frequency interpretation” of probability is shunned by many philosophers and mathematicians; philosophers tend to avoid it because the notion of probability can be given other less mundane metaphysical interpretations, and mathematicians tend to avoid it because they see no need to give the notion of probability any earthly interpretation at all. But in this and most other instances in physical science where the notion of probability has been found to be useful, the frequency interpretation is precisely what is called for. As a bonus, the frequency interpretation allows one to straightforwardly *derive*, not merely postulate, the following three *laws of probability*:

Range law. The probability of an event E is a real number $\Pr(E)$ satisfying $0 \leq \Pr(E) \leq 1$, with the circumstance $\Pr(E) = 0$ corresponding to E *never* occurring, and the circumstance $\Pr(E) = 1$ corresponding to E *always* occurring.

Addition law. If $\Pr(E)$ and $\Pr(F)$ are the respective probabilities of two events E and F , and if these two events are *mutually exclusive* (i.e., they never occur

together), then the probability of the event “either E or F” is $\Pr(E \text{ or } F) = \Pr(E) + \Pr(F)$.

Multiplication law. If $\Pr(E)$ is the probability of an event E, and $\Pr(F|E)$ is the probability of an event F *given* that event E occurs, then the probability of the event “both E and F” is $\Pr(E \ \& \ F) = \Pr(E) \times \Pr(F|E)$.

The “conditional probability of F given E”, which is denoted in the multiplication law by $\Pr(F|E)$, is formally defined in the frequency interpretation to be the ratio of {the number of trials yielding *both* events E and F} to {the number of trials yielding event E *irrespective* of event F}, in the limit of infinitely many trials. The *addition* law evidently requires that $\Pr(F|E) = 0$.

The above three laws of probability will be our *self-evident rules of inference*, not only for our present analysis of the dice tossing problem, but also for our subsequent derivation of the chemical master equation. As for our *self-evident premises* for the dice tossing problem, we shall assume that the phrase “the fair tossing of a fair pair of dice” implies the following:

- (i) Each face of each die has the same probability, namely $\frac{1}{6}$, of turning up in any toss.
- (ii) Knowing the up-face tossed on one die tells us nothing about the up-face tossed on the other die.

Some might object that these two premises are not at all “self-evident”, and that they need to be *proved* by solving appropriate equations of motion for the tossed dice. We might try to counter this objection by noting that experimental confirmation of (i) and (ii) should be easy to obtain. But anyone is free to deny that premises (i) and (ii) are self-evident, and for them the following arguments will *not* constitute a “rigorous” analysis.

With our premises and rules of inference now formulated, let us proceed to deduce the probability of a toss in which the up-face number on die A plus the up-face number on die B equals 4, an event that we shall denote by $(A + B = 4)$. Since the occurrence of this event is really the occurrence of *either* the event $(A = 1 \ \& \ B = 3)$ *or* the event $(A = 2 \ \& \ B = 2)$ *or* the event $(A = 3 \ \& \ B = 1)$, all of which are *mutually exclusive*, then the addition law implies that

$$\begin{aligned} \Pr(A + B = 4) &= \Pr(A = 1 \ \& \ B = 3) + \Pr(A = 2 \ \& \ B = 2) \\ &+ \Pr(A = 3 \ \& \ B = 1) . \end{aligned}$$

The first term on the right here can be written, using the multiplication law, as

$$\Pr(A = 1 \ \& \ B = 3) = \Pr(A = 1) \times \Pr(B = 3 | A = 1) ,$$

where the last factor on the right is the probability that die B turns up 3 given that die A has turned up 1. But since

$$\Pr(A = 1) = \frac{1}{6} \quad (\text{by premise (i)}) ,$$

and

$$\begin{aligned} \Pr(B = 3 \mid A = 1) &= \Pr(B = 3) \quad (\text{by premise (ii)}) \\ &= \frac{1}{6} \quad (\text{by premise (i)}) , \end{aligned}$$

then we have

$$\Pr(A = 1 \ \& \ B = 3) = \frac{1}{6} \times \frac{1}{6} = \frac{1}{36} .$$

Analogous reasoning reveals that the other two terms on the right side of our leading equation are also equal to $\frac{1}{36}$, and so we conclude that

$$\Pr(A + B = 4) = \frac{1}{36} + \frac{1}{36} + \frac{1}{36} = \frac{1}{12} .$$

Have we presented here a “rigorous” analysis of the problem posed? On the negative side, our analysis has none of the usual trappings of a rigor in mathematical physics: we have not written down any Lagrangian or Hamiltonian, and we have not solved any Liouville equation or Schrödinger equation. On the positive side, our analysis would undoubtedly persuade most scientists to lay a very heavy wager on the proposition that the prediction $\frac{1}{12}$ would be confirmed in a carefully performed laboratory experiment. In any case, readers who find that they cannot regard the foregoing analysis as being “rigorous” can spare themselves the trouble of reading the remaining sections, because the analysis there will likewise be found wanting.

3. The chemical system

We consider a system consisting of molecules of N gas-phase chemical species S_1, \dots, S_N , which interreact through M elementary chemical reaction channels R_1, \dots, R_M . We stipulate that the system be kept *well stirred* in a container of constant volume Ω , and *in thermal equilibrium* at a constant absolute temperature T (but not, of course, in *chemical equilibrium*). Precisely what we understand to be implied by these two stipulations will be spelled out shortly. Our ultimate goal will be to describe the behavior of the *species population variables*

$$X_i(t) \equiv \text{the number of molecules of species } S_i \text{ in the system} \\ \text{at time } t \quad (i = 1, \dots, N). \quad (1)$$

The more commonly considered *species concentration variables* are defined by $Z_i(t) \equiv X_i(t)/\Omega$.

We shall focus mainly on elementary reaction channels R_μ of the *bimolecular* form



This form denotes the *collision-initiated* conversion of one molecule of species S_a and one molecule of species S_b into one molecule of species S_c , etc. The two species S_a and S_b are called the R_μ *reactant species*, and they might be the same. There will usually be two R_μ product species S_c and S_d , which also might be the same, but it will not be necessary for us to impose any restrictions here on their number. So-called “reversible” reactions of the form $S_a + S_b \rightleftharpoons S_c + S_d$ should be regarded as two separate elementary reactions. Later we shall comment briefly on when and how *monomolecular* and *trimolecular* reaction channels can be accommodated.

We take the stipulations that the system be “well mixed in a container of volume Ω ” and “in thermal equilibrium at absolute temperature T ” as warrants for the following two premises:

(I) The probability that the center of any randomly selected molecule of the system will be found to lie inside any container subregion of volume $\Delta\Omega$ is equal to $\Delta\Omega/\Omega$.

(II) The probability that the velocity of a randomly selected molecule of mass m will be found to lie inside the infinitesimal region $d^3\mathbf{v}$ about \mathbf{v} is equal to $P_{\text{MB}}(\mathbf{v}) d^3\mathbf{v}$, where

$$P_{\text{MB}}(\mathbf{v}) \equiv \left(\frac{m}{2\pi k_B T} \right)^{3/2} \exp(-m\mathbf{v}^2/2k_B T). \quad (3)$$

Here, \mathbf{v} denotes the Cartesian triplet (v_x, v_y, v_z) , $d^3\mathbf{v} = dv_x dv_y dv_z$, $v^2 = v_x^2 + v_y^2 + v_z^2$, and k_B is Boltzmann’s constant.

Premise (I) asserts that the system is spatially homogeneous, but in a sense that evidently allows for fluctuations from absolute homogeneity. In formal terms, premise (I) asserts that the *position* of a randomly selected molecule can be regarded as a random variable that is uniformly distributed over the interior of the container. Premise (II) is of course the assertion that the molecules have a Maxwell–Boltzmann velocity distribution. In formal terms, premise (II) asserts that each *Cartesian velocity component* of a randomly selected molecule

of mass m can be regarded as a random variable that is normally distributed with mean 0 and variance $k_B T/m$. Implicit in premises (I) and (II) is the assertion that the molecular positions and velocity components are all statistically independent of each other.

We shall not try to decide here whether premises (I) and (II) should be regarded as *theoretically derivable results* or merely *experimentally justifiable postulates*. But in either case, it is clear that these premises derive whatever validity they do have from the extremely sensitive dependence of a classical many-particle system on its initial conditions – a circumstance that ultimately frustrates any *practical* attempt to treat such a system as being “deterministic”. And it is a telling point in favor of premises (I) and (II) that, if they were *not* true, then we probably would *not* be willing to regard the system as being “well mixed and in thermal equilibrium” – just as if premises (i) and (ii) in section 2 were not true then we would not be willing to regard the dice tossing process as being “fair”. In practice, we expect premises (I) and (II) to be valid for any constant-temperature dilute-gas system in which *nonreactive* molecular collisions occur much more frequently than *reactive* molecular collisions.

4. The specific probability rate constant

The first step in our derivation of a “master equation” for the species population variables $\{X_i(t)\}$ will be to establish some pertinent facts about the probability

$$\pi_\mu(t, dt) \equiv \text{probability that a randomly selected combination of } R_\mu \text{ reactant molecules at time } t \text{ will react accordingly in the next infinitesimal time interval } [t, t + dt) \quad (\mu = 1, \dots, M). \quad (4)$$

Specifically, we shall show that for virtually all chemical reaction channels that are consistent with the specifications and assumptions set forth in section 3, $\pi_\mu(t, dt)$ should exist in the form

$$\pi_\mu(t, dt) = c_\mu dt, \quad (5)$$

where c_μ is independent of dt . We call c_μ the *specific probability rate constant* for reaction channel R_μ . In subsection 4.1 we shall calculate $\pi_\mu(t, dt)$ for bimolecular reactions in general, and in subsection 4.2 we shall look at some specific bimolecular mechanisms. In subsection 4.3 we shall consider briefly the matter of monomolecular and trimolecular reactions.

4.1. Bimolecular reactions in general

In order for the bimolecular reaction (2) to occur, it is first necessary that an S_a molecule and an S_b molecule *collide* with each other. Let us assume for simplicity that the molecules of species S_i ($i = a, b$) can be regarded as spheres of mass m_i and radius r_i , so that a collision between an S_a molecule and an S_b molecule occurs whenever the center-to-center distance between those two molecules decreases to $r_a + r_b$. We begin by establishing the existence and mathematical form of the quantity

$$\pi_{\mu}^*(t, dt) \equiv \text{probability that a randomly selected pair of } R_{\mu} \text{ reactant molecules at time } t \text{ will } \textit{collide} \text{ in the next infinitesimal time interval } [t, t + dt). \quad (6)$$

To derive an expression for $\pi_{\mu}^*(t, dt)$, we shall make use of a certain theorem of probability theory. We develop that theorem as follows: Let $E_1, E_2, \dots \equiv \{E_i\}$ be any set of *mutually exclusive and collectively exhaustive* events; i.e., on every trial, one and only one of the events in the set $\{E_i\}$ occurs. Then the occurrence of any event G can be thought of as the occurrence of either the event $E_1 \& G$ or the event $E_2 \& G$ or \dots . Since the events in the set $\{E_i \& G\}$ are *mutually exclusive*, then the addition law implies that

$$\Pr(G) = \sum_i \Pr(E_i \& G).$$

But by the multiplication law we have

$$\Pr(E_i \& G) = \Pr(E_i) \times \Pr(G | E_i).$$

Substituting this into the preceding equation, we get

$$\Pr(G) = \sum_i \Pr(E_i) \times \Pr(G | E_i). \quad (7)$$

If the individual events in the set $\{E_i\}$ happened to be labeled by a *real*-valued index instead of by an integer-valued index, then we must replace the sum in eq. (7) by an *integral* (a sum over nonoverlapping infinitesimal subintervals). We shall in fact be interested in a situation where the integer index i is replaced by a *real 3-dimensional vector* index \mathbf{v} . Eq. (7) then becomes

$$\Pr(G) = \int \int \int_{\mathbf{v}} \Pr(E_{\mathbf{v}}) \times \Pr(G | E_{\mathbf{v}}), \quad (8)$$

valid for $\{E_{\mathbf{v}}\}$ any set of mutually exclusive and collectively exhaustive events.

Now we let G be the event "a randomly selected pair of R_μ reactant molecules at time t collides in the next infinitesimal time interval $[t, t + dt)$ ". Then it follows from definition (6) that

$$\Pr(G) = \pi_\mu^*(t, dt). \quad (9)$$

And we let E_v be the event "a randomly selected pair of R_μ reactant molecules at time t has a relative velocity vector $\mathbf{v}_b - \mathbf{v}_a$ in the infinitesimal Cartesian region $d^3\mathbf{v}$ about \mathbf{v} ". The set of events $\{E_v\}$ is clearly mutually exclusive, and it will also be collectively exhaustive if we allow \mathbf{v} to vary over all of 3-space. As one might expect from elementary kinetic theory,

$$\Pr(E_v) = \left(\frac{m^*}{2\pi k_B T} \right)^{3/2} \exp(-m^*v^2/2k_B T) d^3\mathbf{v}, \quad (10)$$

where v is the magnitude of \mathbf{v} and $m^* \equiv m_a m_b / (m_a + m_b)$ is the reduced mass of the two R_μ reactant molecules. A rigorous proof of eq. (10) can be constructed by using the so-called random variable transformation theorem [2, 3] to calculate the density function of the random variable $\mathbf{V} \equiv \mathbf{V}_b - \mathbf{V}_a$ when each of the two statistically independent random variables \mathbf{V}_a and \mathbf{V}_b has a density function of the form (3).

Finally, we need for eq. (8) an expression for $\Pr(G | E_v)$, the probability that a randomly selected S_a - S_b molecular pair at time t , moving with a relative velocity that is infinitesimally close to \mathbf{v} , will collide with each other in $[t, t + dt)$. In fig. 1 we show a picture of a randomly selected S_a molecule, along with the vector \mathbf{v} that specifies the relative velocity of the otherwise random S_b molecule. It should be evident from the figure that an S_b molecule at time t , moving with velocity \mathbf{v} relative to the S_a molecule, will collide with the S_a molecule in $[t, t + dt)$ if and only if the center of the S_b molecule at time t lies inside the spherically distorted cylindrical region shown in fig. 1 with effective radius $r_a + r_b$ and effective height $v dt$ ^{#1}. So $\Pr(G | E_v)$ is just the probability that the center of a randomly selected S_b molecule, moving with velocity \mathbf{v} relative to the randomly selected S_a molecule, will at time t have its center lying inside a certain subregion of volume $(v dt) \pi(r_a + r_b)^2$. We may therefore conclude from our premise (1) that

$$\Pr(G | E_v) = \frac{(v dt) \pi(r_a + r_b)^2}{\Omega}. \quad (11)$$

^{#1} Here we have implicitly invoked the vanishingly small nature of the infinitesimal dt ; because only if the relative distance $v dt$ is vanishingly small is it permissible to ignore the possibility of collisional interference by some third molecule.




Fig. 1. An S_a molecule of radius r_a , and the infinitesimal region in which the center of an S_b molecule with radius r_b and relative velocity \mathbf{v} must lie if the two molecules are to collide in the next dt . The region is essentially a right cylinder of radius $r_a + r_b$, and height $v dt$, whose upper and lower surfaces have been deformed into identically oriented hemispheres of radii $r_a + r_b$.

Substituting eqs. (9), (10) and (11) into eq. (8), we obtain

$$\pi_{\mu}^*(t, dt) = \int \int \int_{\mathbf{v}} \left(\frac{m^*}{2\pi k_B T} \right)^{3/2} \exp(-m^* v^2 / 2k_B T) d^3 \mathbf{v} \frac{(v dt) \pi (r_a + r_b)^2}{\Omega}. \quad (12)$$

Here, \mathbf{v} is a real 3-dimensional vector variable with unrestricted Cartesian components (v_x, v_y, v_z) , differential element $d^3 \mathbf{v} = dv_x dv_y dv_z$, and magnitude v . Evaluation of this triple integral is straightforward, and yields

$$\pi_{\mu}^*(t, dt) = \Omega^{-1} \left(\frac{8k_B T}{\pi m^*} \right)^{1/2} \pi (r_a + r_b)^2 dt. \quad (13)$$

Of course, we should not suppose that *every* collision between an S_a molecule and an S_b molecule will lead to an R_{μ} reaction. But we shall assume that if such a collision *does* give rise to an R_{μ} reaction, then the reaction happens *immediately*. And we shall define the *collision-conditioned reaction probability* p_{μ} by

$$p_{\mu} \equiv \text{probability that a randomly selected pair of colliding } R_{\mu} \text{ reactant molecules will react according to } R_{\mu}. \quad (14)$$

Since the multiplication law tells us that

$$\text{Pr}(\text{collision \& reaction}) = \text{Pr}(\text{collision}) \times \text{Pr}(\text{reaction} \mid \text{collision}),$$

then by definitions (4), (6) and (14) we have

$$\pi_\mu(t, dt) = \pi_\mu^*(t, dt) p_\mu. \quad (15)$$

Substituting eq. (13) into eq. (15), we conclude that the quantity c_μ defined in eq. (5) exists in the form

$$c_\mu = \Omega^{-1} \left(\frac{8k_B T}{\pi m^*} \right)^{1/2} \pi(r_a + r_b)^2 p_\mu. \quad (16)$$

It will be crucial for our analysis in subsequent sections that c_μ be *independent* of dt . It is clear from our result (16) that this will be the case if and only if p_μ is independent of dt . The definition (14) makes it rather obvious that p_μ will have no dt -dependence; however, it is instructive to check this by examining a few specific bimolecular reaction mechanisms.

4.2. Some simple bimolecular reaction mechanisms

Suppose that each R_μ reactant molecule has a “sensitive area” inscribed on its spherical surface, and that an R_μ reaction will occur if and only if the point of collisional contact between an R_μ reactant pair lies inside both sensitive areas. If the sensitive area on the S_a molecule subtends a solid angle ω_a with respect to the center of the molecule, then owing to the well stirred nature of the system the probability that the collisional contact point will lie inside that area is $\omega_a/4\pi$. Similarly, the probability that the collisional contact point will lie inside the sensitive area on the S_b molecule is $\omega_b/4\pi$. (See fig. 2.) In the absence of any force field that preferentially orients the molecules, the probability that the collisional contact point will simultaneously lie inside both sensitive areas will, by the multiplication law, be the product of these two probabilities. Therefore, the collision-conditioned reaction probability in this case is simply

$$p_\mu = \omega_a \omega_b / (4\pi)^2. \quad (17)$$

Another kind of bimolecular reaction model concerns the *impact energy*, which we define, following Present [4], by

$$\varepsilon \equiv \text{the kinetic energy associated with the component of the relative velocity of the colliding molecules along their line of centers at contact.} \quad (18)$$




Fig. 2. A collision between an S_a molecule and an S_b molecule in which the point of collisional contact lies *inside* a certain solid angle ω_a on the S_a molecule and *outside* a certain solid angle ω_b on the S_b molecule.

This variable would typically influence a reaction R_μ through some function H_μ according to

$$H_\mu(\varepsilon) \equiv \text{probability that a pair of } R_\mu \text{ reactant molecules, colliding with impact energy } \varepsilon, \text{ will undergo an } R_\mu \text{ reaction.} \quad (19)$$

To derive an expression for the collision-conditioned reaction probability p_μ in this case, it is necessary to make a direct calculation of the reaction probability $\pi_\mu(t, dt)$ defined in (4); we can then deduce p_μ from eqs. (15) and (13). For this calculation, we shall require an easily proved extension of theorem (8): If G is any event, $\{E_\nu\}$ is any set of mutually exclusive and collectively exhaustive events labeled by the real vector index ν , and $\{F_\beta\}$ is any set of mutually exclusive and collectively exhaustive events labeled by the real scalar index β , then

$$\Pr(G) = \int \int \int \int_{\beta} \Pr(E_\nu) \times \Pr(F_\beta | E_\nu) \times \Pr(G | E_\nu \& F_\beta). \quad (20)$$

We let G be the event “a randomly selected pair of R_μ reactant molecules at time t reacts accordingly in the next infinitesimal time interval $[t, t + dt)$ ”. Definition (6) then gives

$$\Pr(G) = \pi_\mu(t, dt). \quad (21)$$

We let E_v be the event "a randomly selected pair of R_μ reactant molecules at time t has a relative velocity vector $\mathbf{v}_b - \mathbf{v}_a$ in the infinitesimal Cartesian region $d^3\mathbf{v}$ about \mathbf{v} ". This is the same definition for E_v that was adopted in subsection 4.1, so $\Pr(E_v)$ is as given in eq. (10).

We let F_β be the event "a randomly selected pair of R_μ reactant molecules at time t collides in the next infinitesimal time interval $[t, t + dt)$ with impact parameter in the infinitesimal interval $[\beta, \beta + d\beta)$ ". The impact parameter is defined to be the perpendicular distance between the precollision trajectories of the colliding molecules; so if β is allowed to range from 0 to $r_a + r_b$ then the set of events $\{F_\beta\}$ will indeed be mutually exclusive and collectively exhaustive. For eq. (20), we require the probability of the event F_β given the event E_v , i.e., given that the colliding molecules have at time t a relative velocity vector $\mathbf{v}_b - \mathbf{v}_a$ that is infinitesimally close to the vector \mathbf{v} . In fig. 3 we show a picture of a randomly selected S_a molecule, along with the vector \mathbf{v} that specifies the relative velocity of the otherwise random S_b molecule. It should be clear from the figure that an S_b molecule at time t , moving with velocity \mathbf{v} relative to the S_a molecule, will collide with the S_a molecule in $[t, t + dt)$ with impact parameter in $[\beta, \beta + d\beta)$ if and only if the center of the S_b molecule at time t lies in the cylindrical shell shown in fig. 3 of radius β , wall thickness $d\beta$, and height $v dt$ (see footnote 1). So $\Pr(F_\beta | E_v)$ is the probability that the center of




Fig. 3. An S_a molecule, and the infinitesimal region in which the center of an S_b molecule with relative velocity \mathbf{v} must lie if the two molecules are to collide in the next dt with impact parameter between β and $\beta + d\beta$. The region is essentially a cylindrical shell of inner radius β , outer radius $\beta + d\beta$ and height $v dt$, with the upper and lower surfaces conformed upon a sphere of radius $r_a + r_b$.

a randomly selected S_b molecule, moving with velocity \mathbf{v} relative to the randomly selected S_a molecule, will at time t have its center lying inside a certain subregion of volume $(v dt)(2\pi\beta)d\beta$. It follows from premise I that

$$\Pr(F_\beta | E_v) = \frac{(v dt)(2\pi\beta d\beta)}{\Omega}. \quad (22)$$

Finally, we need for eq. (20) an expression for $\Pr(G | E_v \& F_\beta)$, the probability that a randomly selected S_a - S_b molecular pair at time t , moving with a relative velocity that is infinitesimally close to \mathbf{v} and colliding in $[t, t + dt)$ with an impact parameter that is infinitesimally close to β , will actually undergo an R_μ reaction. According to our hypothesis (19), this probability is just $H_\mu(\varepsilon)$ where ε is the impact energy defined in (18). Fig. 4 shows the geometry of the colliding molecules, and it is clear from this figure that ε is equal to $(m^*/2)(v \cos \theta)^2$, where θ is the angle between \mathbf{v} and the line of centers at contact. The figure also shows that $\sin \theta = \beta / (r_a + r_b)$. Thus we conclude that

$$\Pr(G | E_v \& F_\beta) = H_\mu(\varepsilon(v, \beta)), \quad (23)$$

where $\varepsilon(v, \beta)$ is the function

$$\varepsilon(v, \beta) \equiv \frac{m^* v^2}{2} \left[1 - \left(\frac{\beta}{r_a + r_b} \right)^2 \right]. \quad (24)$$




Fig. 4. The geometry of an S_a - S_b molecular collision with impact parameter β .

Substituting into eq. (20) the expressions for the various probabilities in eqs. (21), (10), (22) and (23), we obtain

$$\pi_{\mu}(t, dt) = \int \int \int \int \left(\frac{m^*}{2\pi k_B T} \right)^{3/2} \exp(-m^* v^2 / 2k_B T) d^3 \mathbf{v} \\ \times \frac{(v dt)(2\pi\beta d\beta)}{\Omega} H_{\mu}(\varepsilon(v, \beta)). \quad (25)$$

In this formula, \mathbf{v} is a vector variable with unrestricted Cartesian components (v_x, v_y, v_z) , differential element $d^3 \mathbf{v} = dv_x dv_y dv_z$, and magnitude v ; β is a scalar variable ranging from 0 to $r_a + r_b$; H_{μ} is the function defined in eq. (19); and $\varepsilon(v, \beta)$ is the function defined in eq. (24).

An evaluation of the four-fold integral (25) obviously cannot be carried out unless the function H_{μ} is specified. As an example, suppose that $H_{\mu}(\varepsilon)$ is unity for $\varepsilon > \varepsilon_{\mu}$ and zero for $\varepsilon \leq \varepsilon_{\mu}$, where ε_{μ} is some given *threshold impact energy*. In that case, eq. (25) would read

$$\pi_{\mu}(t, dt) = \int_{-\infty}^{\infty} dv_x \int_{-\infty}^{\infty} dv_y \int_{-\infty}^{\infty} dv_z \int_0^{r_a+r_b} d\beta \left(\frac{m^*}{2\pi k_B T} \right)^{3/2} \\ \times \exp(-m^* v^2 / 2k_B T) \frac{(v dt)(2\pi\beta)}{\Omega} \Theta(\varepsilon(v, \beta) - \varepsilon_{\mu}), \quad (26a)$$

where $\Theta(x)$ is unity for $x > 0$ and zero otherwise. Evaluating this integral is straightforward, and when the result is combined with eqs. (15) and (13) we find that p_{μ} has the ‘‘Arrhenius form’’

$$p_{\mu} = \exp(-\varepsilon_{\mu} / k_B T). \quad (26b)$$

It seems clear from the foregoing simple examples of bimolecular reaction mechanisms that the factor p_{μ} in the c_{μ} formula (16) should always exist and be independent of dt .

4.3. Monomolecular and trimolecular reactions

A *monomolecular* reaction has the form



and describes the spontaneous conversion of an S_a molecule into an S_c molecule. This reaction should not be confused with the *bimolecular* reaction



in which the conversion of S_a to S_c is “catalyzed” by a collision with an S_b molecule. We must presume that the monomolecular reaction (27) is caused by some internal quantum mechanical mechanism analogous to that responsible for the decay of a radioactive nucleus. The time-dependent perturbation theory of quantum mechanics shows that most weakly perturbed energy-conserving transitions have occurrence probabilities in times δt that are linear in δt when δt is *microscopically* large but *macroscopically* small; i.e., to a good approximation the probability for a radioactive nucleus to decay in the next dt has the mathematical form αdt , where α is some constant. On the basis of this analogy, we may expect $\pi_\mu(t, dt)$ for the monomolecular reaction (27) to exist in the *approximate* form $c_\mu dt$, where c_μ is independent of dt .

As for *trimolecular* reactions of the form



a good argument can be made that they do not really occur as “elementary events”, and therefore should not be considered. There is, however, one instance in which a set of three coupled reaction channels, with a total of four reactant molecules, can be *approximately* described by a single trimolecular reaction channel whose reaction probability $\pi_\mu(t, dt)$ is effectively linear in dt . Just how this comes about is discussed on pp. 359–361 of ref. [3]. So if the conditions for validity of this approximate description are adequately satisfied, then it should be permissible to invoke a trimolecular reaction channel for which $\pi_\mu(t, dt)$ *approximately* exists in the form $c_\mu dt$, where c_μ is independent of dt .

5. The chemical master equation

In section 4 we endeavored to establish the following basic fact: For each of the M elementary reaction channels R_μ ($\mu = 1, \dots, M$) open to the molecules of a well stirred, thermally equilibrated gas-phase system, there exists a dt -independent scalar c_μ such that

$$c_\mu dt = \text{probability that a randomly selected combination of } R_\mu \text{ reactant molecules at time } t \text{ will react accordingly in the next infinitesimal time interval } [t, t + dt). \quad (30)$$


Now, the specific probability rate constant c_μ is but one of *three* entities that we

shall require to fully characterize a particular reaction channel R_μ . The other two entities are a function h_μ and a set of N integers $(\nu_{\mu 1}, \dots, \nu_{\mu N})$, and they are defined as follows:

$$h_\mu(n_1, \dots, n_N) \equiv \text{the number of distinct combinations of } R_\mu \text{ reactant molecules in the system when there are exactly } n_i \text{ of the } S_i \text{ molecules } (i = 1, \dots, N), \quad (31)$$

$$\nu_{\mu i} \equiv \text{the change in the } S_i \text{ molecular population caused by the occurrence of one } R_\mu \text{ reaction.} \quad (32)$$

The function h_μ and the integers $\{\nu_{\mu i}\}$ can be straightforwardly deduced by inspecting the ‘‘algebraic structure’’ of reaction channel R_μ . For example, consider the reaction channels



The h_μ functions for these channels are

$$h_1(n_1, \dots, n_N) = n_1 n_2, \quad (34a)$$

$$h_2(n_1, \dots, n_N) = n_1(n_1 - 1)/2, \quad (34b)$$

$$h_3(n_1, \dots, n_N) = n_1. \quad (34c)$$

And the $\nu_{\mu i}$ values for these channels are

$$\nu_{11} = -1, \quad \nu_{12} = -1, \quad \nu_{13} = +1, \quad \nu_{14} = +1, \quad \text{all other } \nu_{1i} = 0, \quad (35a)$$

$$\nu_{21} = -1, \quad \nu_{22} = +1, \quad \text{all other } \nu_{2i} = 0, \quad (35b)$$

$$\nu_{31} = -1, \quad \nu_{32} = +1, \quad \text{all other } \nu_{3i} = 0. \quad (35c)$$

Notice that the functional form of h_μ uniquely specifies the ‘‘reactant side’’ of R_μ , and the values of $\nu_{\mu 1}, \dots, \nu_{\mu N}$ then uniquely specify the ‘‘product side’’ of R_μ .

To simplify the writing in what follows, we shall adopt the vector notation

$$\mathbf{n} \equiv (n_1, \dots, n_N),$$

where the components n_i ($i = 1, \dots, N$) of \mathbf{n} are understood to be *integer-valued*. In particular, we shall write $X(t) \equiv (X_1(t), \dots, X_N(t))$ and $\mathbf{v}_\mu \equiv (v_{\mu 1}, \dots, v_{\mu N})$. Notice that if the components of vector \mathbf{n} give the molecular populations of the various species immediately *before* the occurrence of an R_μ reaction, then the components of the vector $\mathbf{n} + \mathbf{v}_\mu$ give the molecular populations immediately *after* the occurrence of that reaction.

Our task is to describe the behavior of the species population vector $X(t)$ under the influence of the M reaction channels R_μ ($\mu = 1, \dots, M$), where each R_μ is specified by its specific probability rate constant c_μ , its reactant combination function h_μ , and its jump vector \mathbf{v}_μ . To accomplish this task, we begin by proving three theorems.

Theorem 1. If $X(t) = \mathbf{n}$, then the probability that *exactly one* R_μ reaction will occur in the system in the time interval $[t, t + dt)$ is equal to $c_\mu h_\mu(\mathbf{n}) dt + o(dt)$, where $o(dt)$ denotes terms that go with zero with dt faster than dt .

Proof. Imagine that every molecule of each reactant species in the system at time t is assigned a unique label. Since the system is “well stirred”, then calling out the labels of a particular combination of R_μ reactant molecules is equivalent to “randomly selecting such a combination. So by eq. (30), each of the $h_\mu(\mathbf{n})$ distinct combinations of R_μ reactant molecules in the system at time t has probability $c_\mu dt$ of reacting according to R_μ in $[t, t + dt)$. And by the range and addition laws, each has probability $1 - c_\mu dt$ of *not* reacting thus in $[t, t + dt)$. The multiplication law then implies that the probability that a *particular one* of the $h_\mu(\mathbf{n})$ R_μ reactant combinations *does* react according to R_μ in $[t, t + dt)$ while the other $h_\mu(\mathbf{n}) - 1$ combinations do *not*, is

$$c_\mu dt (1 - c_\mu dt)^{h_\mu(\mathbf{n}) - 1} = c_\mu dt + o(dt) .$$

We can now use the addition law to calculate the probability that *any* of the $h_\mu(\mathbf{n})$ distinct R_μ reactant combinations at time t will react *alone* in $[t, t + dt)$ as the sum of their separate probabilities (because these events are mutually exclusive). Since each of the $h_\mu(\mathbf{n})$ terms in this probability sum is equal to $c_\mu dt + o(dt)$, then the sum is equal to

$$h_\mu(\mathbf{n}) [c_\mu dt + o(dt)] = c_\mu h_\mu(\mathbf{n}) dt + o(dt) .$$

QED

Theorem 2. If $X(t) = \mathbf{n}$, then the probability that *no* reaction will occur in the system in the time interval $[t, t + dt)$ is equal to $1 - \sum_{\mu} c_{\mu} h_{\mu}(\mathbf{n}) dt + o(dt)$.

Proof. As was noted in the proof of theorem 1, each of the $h_{\mu}(\mathbf{n})$ distinct combinations of R_{μ} reactant molecules in the system at time t has probability $1 - c_{\mu} dt$ of *not* reacting according to R_{μ} in $[t, t + dt)$. So by the multiplication law, the probability that *all* of those $h_{\mu}(\mathbf{n})$ combinations of R_{μ} reactant molecules will *not* react according to R_{μ} in $[t, t + dt)$, is

$$(1 - c_{\mu} dt)^{h_{\mu}(\mathbf{n})} = 1 - h_{\mu}(\mathbf{n}) c_{\mu} dt + o(dt) .$$

Therefore, appealing once again to the multiplication law, the probability that no R_1 reaction *and* no R_2 reaction . . . *and* no R_M reaction will occur in $[t, t + dt)$, is

$$\prod_{\mu=1}^M [1 - h_{\mu}(\mathbf{n}) c_{\mu} dt + o(dt)] = 1 - \sum_{\mu=1}^M h_{\mu}(\mathbf{n}) c_{\mu} dt + o(dt) . \quad \text{QED}$$

Theorem 3. The probability of more than one reaction occurring in the system in the time interval $[t, t + dt)$ is $o(dt)$.

Proof. From the multiplication law and eq. (30), we see that the probability for k reactions to occur in $[t, t + dt)$ must be proportional to $(dt)^k$, and therefore will be $o(dt)$ for $k > 1$. QED

Armed with the foregoing three theorems, we are now in a position to develop an analytical description of the behavior of the species population vector $X(t)$. It should be obvious that even if we fix the state of the system at some initial time t_0 ,

$$X(t_0) = \mathbf{n}_0 , \quad (36)$$

we cannot hope to develop a deterministic time-evolution equation for $X(t)$ for $t > t_0$. So we shall try instead to derive a deterministic time-evolution equation for the *probability* function

$$P(\mathbf{n}, t | \mathbf{n}_0, t_0) = \text{probability that } X(t) = \mathbf{n}, \text{ given that } X(t_0) = \mathbf{n}_0 \\ (t \geq t_0) . \quad (37)$$

Our strategy for constructing a time-evolution equation for this function P is to use our three theorems, along with the addition and multiplication laws of

probability, to express $P(\mathbf{n}, t + dt | \mathbf{n}_0, t_0)$ as the sum of the probabilities of some mutually exclusive and collectively exhaustive routes from $X(t_0) = \mathbf{n}_0$ to $X(t + dt) = \mathbf{n}$, these routes being distinguished from one another by *what happens in the time interval* $[t, t + dt)$. In the following paragraph, we shall prove by such reasoning that the probability $P(\mathbf{n}, t + dt | \mathbf{n}_0, t_0)$ can be written

$$\begin{aligned}
 P(\mathbf{n}, t + dt | \mathbf{n}_0, t_0) &= P(\mathbf{n}, t | \mathbf{n}_0, t_0) \times \left(1 - \sum_{\mu=1}^M c_{\mu} h_{\mu}(\mathbf{n}) dt + o(dt) \right) \\
 &+ \sum_{\mu=1}^M P(\mathbf{n} - \mathbf{v}_{\mu}, t | \mathbf{n}_0, t_0) \times [c_{\mu} h_{\mu}(\mathbf{n} - \mathbf{v}_{\mu}) dt + o(dt)] \\
 &+ o(dt). \tag{38}
 \end{aligned}$$

One possible route from $X(t_0) = \mathbf{n}_0$ to $X(t + dt) = \mathbf{n}$ is for *no reaction* to occur in $[t, t + dt)$. For this route, the system obviously must be in state \mathbf{n} at time t . So it follows from the multiplication law that the probability of this “no reaction in $[t, t + dt)$ ” route is equal to the product of {the probability that the system will go from $X(t_0) = \mathbf{n}_0$ to $X(t) = \mathbf{n}$ } times {the subsequent probability that no reaction will occur in $[t, t + dt)$ }. In light of the definition (37) and theorem 2, this product is just the first term on the right-hand side of eq. (38). Another possible route from $X(t_0) = \mathbf{n}_0$ to $X(t + dt) = \mathbf{n}$ is for *exactly one* R_{μ} reaction to occur in $[t, t + dt)$. For this route, the system obviously must be in state $\mathbf{n} - \mathbf{v}_{\mu}$ at time t . So it follows from the multiplication law that the probability of this “one R_{μ} reaction in $[t, t + dt)$ ” route is equal to the product of {the probability that the system will go from $X(t_0) = \mathbf{n}_0$ to $X(t) = \mathbf{n} - \mathbf{v}_{\mu}$ } times {the subsequent probability that exactly one R_{μ} reaction will occur in $[t, t + dt)$ }. In light of the definition (37) and theorem 1, this product is just the μ th summand of the second term on the right-hand side of eq. (38). (There will of course be M routes of this kind, one for each reaction channel R_{μ} .) Finally, if the system’s journey from $X(t_0) = \mathbf{n}_0$ to $X(t + dt) = \mathbf{n}$ is neither by the “no reaction in $[t, t + dt)$ ” route nor by any of the “one reaction in $[t, t + dt)$ ” routes, then it obviously must be by a “more than one reaction in $[t, t + dt)$ ” route. According to theorem 3, the probability of that kind of route is $o(dt)$. Now, all these routes are *mutually exclusive*, because they have mutually exclusive things taking place in $[t, t + dt)$; so the addition law allows us to calculate the probability that *any* of them will occur by simply summing their separate probabilities. This is precisely the statement that is made by eq. (38).

We now subtract $P(\mathbf{n}, t | \mathbf{n}_0, t_0)$ from both sides of eq. (38), divide through by dt , and then take the limit $dt \downarrow 0$. Since all $o(dt)/dt$ terms vanish in this limit, we evidently obtain

$$\frac{\partial}{\partial t} P(\mathbf{n}, t | \mathbf{n}_0, t_0) = \sum_{\mu=1}^M [c_{\mu} h_{\mu}(\mathbf{n} - \mathbf{v}_{\mu}) P(\mathbf{n} - \mathbf{v}_{\mu}, t | \mathbf{n}_0, t_0) - c_{\mu} h_{\mu}(\mathbf{n}) P(\mathbf{n}, t | \mathbf{n}_0, t_0)]. \quad (39)$$

This is the *chemical master equation*. It is a t -evolution equation for the function $P(\mathbf{n}, t | \mathbf{n}_0, t_0)$, for fixed \mathbf{n}_0 and t_0 , and it is to be solved subject to the initial condition

$$P(\mathbf{n}, t = t_0 | \mathbf{n}_0, t_0) = \begin{cases} 1, & \text{if } \mathbf{n} = \mathbf{n}_0, \\ 0, & \text{if } \mathbf{n} \neq \mathbf{n}_0, \end{cases} \quad (40)$$

as required by eq. (36).

6. Discussion

The foregoing derivation of the chemical master equation is a greatly refined version of arguments advanced earlier by the author [5] in developing the so-called *stochastic simulation algorithm*. It may therefore be appropriate to clarify here precisely how the stochastic simulation algorithm fits into the theoretical framework just elaborated.

The key to the stochastic simulation algorithm is not the probability function $P(\mathbf{n}, t | \mathbf{n}_0, t_0)$ that appears in the chemical master equation, but rather another probability function $p(\tau, \mu | \mathbf{n}, t)$, which is defined as follows:

$$p(\tau, \mu | \mathbf{n}, t) d\tau \equiv \text{probability that, given } X(t) = \mathbf{n}, \text{ the next reaction in the system will occur in the infinitesimal time interval } [t + \tau, t + \tau + d\tau), \text{ and will be an } R_{\mu} \text{ reaction.} \quad (41)$$

In the language of random variable theory, p is the joint density function for the two random variables "time to the next reaction" and "index of the next reaction", with the possible values of these two random variables being represented respectively by the real variable τ ($0 \leq \tau < \infty$) and the integer variable μ ($1 \leq \mu \leq M$). To derive an explicit formula for p , let

$$a(\mathbf{n}) \equiv \sum_{\mu=1}^M c_{\mu} h_{\mu}(\mathbf{n}), \quad (42)$$

and imagine the time interval $[t, t + \tau + d\tau)$ to be subdivided into $k + 1$ subintervals, where $k > 1$. As indicated in fig. 5, the first k of these subintervals are taken to be of equal length $\varepsilon = \tau/k$, and together they cover the interval




Fig. 5. A partitioning of the time interval $[t, t + \tau + d\tau]$ into $k + 1$ nonoverlapping subintervals.

$[t, t + \tau]$; the last subinterval is the interval $[t + \tau, t + \tau + d\tau]$. With $X(t) = \mathbf{n}$, $p(\tau, \mu | \mathbf{n}, t) d\tau$ is evidently the probability of the event “no reaction occurs in each of the k ε -subintervals and exactly one R_μ reaction occurs in the final $d\tau$ -subinterval”. So, recalling theorems 1 and 2 of section 5, and taking account of the definition (42), we have by the multiplication law that

$$p(\tau, \mu | \mathbf{n}, t) d\tau = [1 - a(\mathbf{n}) \varepsilon + o(\varepsilon)]^k [c_\mu h_\mu(\mathbf{n}) d\tau + o(d\tau)].$$

Dividing this equation through by $d\tau$ and taking the limit $d\tau \downarrow 0$, we get

$$p(\tau, \mu | \mathbf{n}, t) = [1 - a(\mathbf{n}) \varepsilon + o(\varepsilon)]^k c_\mu h_\mu(\mathbf{n}). \quad (43)$$

Eq. (43) is valid for *any* integer $k > 1$, and so it is valid in the limit $k \rightarrow \infty$. Before taking that limit, let us write the first factor on the right as

$$\begin{aligned} [1 - a(\mathbf{n}) \varepsilon + o(\varepsilon)]^k &= \left(1 - \frac{a(\mathbf{n}) k\varepsilon + k o(\varepsilon)}{k}\right)^k \\ &= \left(1 - \frac{a(\mathbf{n}) \tau + \tau[o(\varepsilon)/\varepsilon]}{k}\right)^k, \end{aligned}$$

where the last step uses the fact that $k\varepsilon = \tau$. Now letting $k \rightarrow \infty$, and noting that $o(\varepsilon)/\varepsilon \downarrow 0$ in that limit, we get

$$\lim_{k \rightarrow \infty} [1 - a(\mathbf{n}) \varepsilon + o(\varepsilon)]^k = \lim_{k \rightarrow \infty} \left(1 - \frac{a(\mathbf{n}) \tau}{k}\right)^k = e^{-a(\mathbf{n}) \tau}.$$

Substituting this result into eq. (43), and then multiplying and dividing by $a(\mathbf{n})$, we conclude that the function p defined in eq. (41) is given by

$$p(\tau, \mu | \mathbf{n}, t) = a(\mathbf{n}) e^{-a(\mathbf{n}) \tau} \frac{c_\mu h_\mu(\mathbf{n})}{a(\mathbf{n})}. \quad (44)$$

Eq. (44) provides the mathematical basis for the stochastic simulation algorithm. Without going into all the details here, eq. (44) says that, given $X(t) = \mathbf{n}$, the “time to the next reaction” and the “index of the next reaction”

are statistically independent random variables, the former having the exponential density function with decay constant $a(\mathbf{n})$, and the latter having the integer density function $c_\mu h_\mu(\mathbf{n})/a(\mathbf{n})$. It can be shown [3, 5] that we can generate a statistically exact sample pair (τ, μ) according to this joint density function by first generating two unit-interval uniform random numbers r_1 and r_2 , then taking τ to be

$$\tau = [1/a(\mathbf{n})] \ln(1/r_1), \quad (45a)$$

and finally taking μ to be the *smallest* integer for which

$$\sum_{\mu'=1}^{\mu} c_{\mu'} h_{\mu'}(\mathbf{n}) > r_2 a(\mathbf{n}). \quad (45b)$$

With values for τ and μ chosen thusly, we then “advance the system’s state vector” from $X(t) = \mathbf{n}$ to $X(t + \tau) = \mathbf{n} + \mathbf{v}_\mu$. Repeated application of this advancement procedure is the essence of the stochastic simulation algorithm. It should be emphasized that this advancement procedure is *exact* to the extent that r_1 and r_2 are “fair samplings” of the unit-interval uniform random variable. In particular, the advancement from t to $t + \tau$ is *not* a finite approximation of an infinitesimal time step, as in a standard differential equation solver; instead, the species population vector maintains the value \mathbf{n} throughout the entire finite time interval $[t, t + \tau)$, and then abruptly changes to $\mathbf{n} + \mathbf{v}_\mu$ at the instant $t + \tau$ when the R_μ reaction occurs.

Since the chemical master equation and the stochastic simulation algorithm are derived from the same set of theorems, then they are *logically equivalent* to each other. In more precise terms, the stochastic simulation algorithm produces exact “realizations” of the jump Markov process $X(t)$ whose initially conditioned density function is determined by the chemical master equation. In essence, we have proved here that both the chemical master equation and the stochastic simulation algorithm are rigorous consequences of premises (I) and (II).

References

- [1] D.A. McQuarrie, J. Appl. Prob. 4 (1967) 413.
- [2] D.T. Gillespie, Am. J. Phys. 51 (1983) 520.
- [3] D.T. Gillespie, Markov Processes: An Introduction for Physical Scientists (Academic Press, San Diego, 1992).
- [4] R.D. Present, Kinetic Theory of Gases (McGraw-Hill, New York, 1958) ch. 8.
- [5] D.T. Gillespie, J. Comput. Phys. 22 (1976) 403; J. Phys. Chem. 81 (1977) 2340.