

Inferring and Propagating Kinetic Parameter Uncertainty for Condensed Phase Burning Models

Morgan C. Bruns

Received: date / Accepted: date

Abstract Kinetic parameters for serial pyrolysis reactions were calibrated from thermogravimetric analysis (TGA) data using Bayesian inference via Markov Chain Monte Carlo (MCMC) simulations assuming a serial reaction mechanism. Calibrations were performed for high-impact polystyrene (HIPS), bisphenol-A polycarbonate (PC), and poly(vinyl chloride) (PVC) at heating rates of 3 K/min and 10 K/min. The resulting parameter inferences are probabilistic as opposed to the point estimates calibrated in previous studies and are visualized using posterior probability density functions (PDFs) generated by kernel density estimation (KDE). Correlations between the parameters are identified and discussed. In particular, it is clear that pre-exponential constants and activation energies for a given reaction have a strong positive correlation. It is hypothesized that the degree of overlap in the posterior PDFs might be a measure of model adequacy. Point-estimates of the kinetic parameters were made by finding the mode of the posterior PDFs. For HIPS, it was determined that a one-reaction pyrolysis model is most appropriate, and that the posterior modes for $\log(A_1)$ and E_1 are $19.5 \log(1/s)$ and 292 kJ/mol, respectively, for the 3 K/min data. To evaluate the effect of kinetic parameter uncertainty on predictions of burning rate, samples from the posterior PDF were used to simulate gasification and cone calorimetry experiments using the Fire Dynamics Simulator (FDS). In some cases, it was found that models with fewer parameters provided better predictions due to over-fitting associated with greater model complexity. Another important observation is that for the predictions of PVC cone calorimetry, the time to peak heat release rate can range from around 40 s to 180 s for a number of different kinetic parameter combinations that all fit the TGA data fairly well. It is argued that the proposed methodology is necessary for progress in modeling of condensed phase physics for fire problems as it supports both model validation and engineering predictions.

Keywords Pyrolysis · Kinetics · TGA · Bayesian Inference · Markov Chain Monte Carlo · Uncertainty Propagation · FDS

M. Bruns
National Institute of Standards and Technology, 100 Bureau Drive, Mail Stop 8665, Gaithersburg, MD 20899, USA
Tel.: +301-975-2088
Fax: +301-975-4052
E-mail: morgan.bruns@nist.gov

1 Introduction

Computational tools for predicting various aspects of compartment and building fires have found increasing use in engineering practice. The development of these tools has focused primarily on modeling gas phase physical processes such as turbulence and combustion. In particular, the Fire Dynamics Simulator (FDS) has been validated for a variety of scenarios in which gas phase processes dominate such as compartment temperatures, plume temperatures, and heat fluxes [1]. However, models for heat and mass transfer in the condensed phase of burning materials have been much less studied. Because of this, computational predictions of burning rate and flame spread involve a large amount of uncertainty for common combustible objects. Since these processes are important for predicting the development of building fires, it would be of considerable value to reduce this uncertainty by improving condensed phase modeling capabilities.

Several models for predicting condensed phase burning have been developed including the FDS solid phase submodel [2], Gpyro [3], and ThermaKin [4]. For brevity, such condensed phase models will be referred to as burning models in the following. These models must overcome several challenges if burning is to become computationally predictable. One of the principle difficulties faced by burning models is the lack of direct measurement techniques for estimating model parameters. Condensed phase matter is inherently more complex than gases, and so much of the experimental techniques and theory that have proven successful for the gas phase are not applicable for solids and liquids. For thermophysical properties such as heat capacity and thermal conductivity, the measurement problem is compounded by the fact that detailed models require values for these properties as a function of both temperature and composition—both of which change significantly during burning. In the experiments that do exist, it is not easy to separate the effects of the many relevant physical properties on the measured quantities.

The process of determining the values of model parameters from experimental data is referred to as calibration. For condensed phase burning models, many authors have made use of optimization algorithms to calibrate the many material properties parameterizing their models [5,3,6]. Some work has been done to determine which optimization algorithms are most effective for this class of problems. Lautenberger and Fernandez-Pello [7] compared four algorithms for finding 19 model parameters that best fit synthetic data generated by a 1d condensed phase burning model. It was found that the shuffled complex evolution algorithm (SCE), which was first applied to burning model parameter estimation by Chaos et al. [6], consistently and efficiently found the true parameters used to generate the synthetic data. However, several problems exist with optimization-based approaches to model calibration. First, there is no known test for determining if the global optimum has been found. Second, the only output of optimization algorithms is a point estimate of the calibrated material properties—i.e., there is no characterization of the uncertainty in these parameters as implied by the data. Third, the papers referenced above calibrate against data from physically complex experiments such as gasification, cone calorimetry, or the fire propagation apparatus. In such scenarios, many physical processes influence the measured quantities. Consequently, confounding effects typically arise such that the calibrated parameters assume unrealistic values, and extrapolation to scenarios different from the calibration experiments is unjustifiable. It is therefore advisable to use experiments in which only a small number of parameters influence the measured output. In this paper, thermogravimetric analysis (TGA) experiments are used for calibration since the measured output depends primarily on the kinetic parameters.

Many methods for calibrating thermal decomposition kinetic parameters have been presented in the literature [8,9]. However, all of these methods are deterministic in the sense of producing point estimates for the parameter values. Point estimates can be misleading in that model predictions based on them might be significantly less certain than the TGA data actually implies. A more robust approach is founded on Bayesian data analysis [10, 11]. The Bayesian approach to model calibration treats the model parameters as quantities whose values are imprecisely known. These uncertain model parameters are characterized by probability density functions (PDFs) instead of point estimates. The form of these PDFs are determined by a calibration process in which Bayes' theorem is used to update prior knowledge with experimental data. The value of Bayesian calibration is summarized in a recent report on predictive computational science [12]:

We prefer the Bayesian approach over traditional deterministic calibration for two primary reasons. First, unlike deterministic estimates, the solution of the Bayesian inverse problem is a complete PDF describing what parameter values are consistent with prior information and experimental observations. Second, the Bayesian formulation requires a complete treatment of uncertainty. This treatment is essential in making correct inferences from the data.

Bayesian inference has been applied to fire problems by Overholt and Ezekoye [13]. This article is recommended as it contains a more thorough overview of the theory and the computational framework used in the present paper. A similar study to the present paper was presented in [14]. In that work, a Bayesian calibration was performed against gasification data similar to that used in Section 4 of this paper. In contrast, the present paper studies the feasibility and value of using small scale experiments to perform a Bayesian calibration and then using the results of the small scale calibration to make predictions regarding more complicated condensed phase burning scenarios. There are many chemical and thermophysical parameters influencing burning rate. Some of these parameters will have a significant impact on the model predictions. Since measurements and calibrations can be expensive and time-consuming, it is advisable to focus only on those parameters to which the model predictions are most sensitive. Detailed sensitivity analyses of condensed phase burning models have been performed in the literature [15, 16]. Both of these studies found that model predictions assuming nominal polymer property values were either most or second most sensitive to changes in the pyrolysis kinetic parameters. For this reason, it is of considerable practical value to characterize the uncertainty in kinetic parameters in order to understand the uncertainty in model predictions.

A related topic to model calibration is model selection. Logically, model selection precedes calibration since selection determines the parameters that must be calibrated. In practice, some degree of iteration between selection and calibration is typically required. Roughly speaking, a model should incorporate as much physics as is necessary for predicting the quantity of interest. Including more physics introduces more parameters that must be measured or calibrated from experimental data. Therefore, it is not reasonable to include unnecessary physics. Bal and Rein [17] studied model selection in the context of the thermal degradation of non-charring polymers. By comparing simulation results with data for several different models, it was found that complex chemical mechanisms are not necessarily required to accurately predict surface temperatures and mass loss rate in thermally degrading poly(methyl methacrylate) (PMMA). Bayesian inference provides a natural framework for model selection since Bayes' theorem provides the rule for determining the relative probability of a model from available experimental data. Several approaches have been suggested for doing Bayesian model selection including the Bayesian information criterion (BIC) [18]

and the deviance information criterion (DIC) [10]. Such approaches are a natural extension of the techniques used in this paper as they make use of the likelihood function discussed below.

In this study, a more thorough approach to Bayesian calibration, validation, and prediction is demonstrated. First, it is shown that Bayesian inference can be used for thermal degradation kinetic parameter calibration. Correlations between various kinetic parameters are observed and studied. Second, the effect of the implied kinetic parameter uncertainty on solid burning rate simulations is examined. These results demonstrate how model validation is impacted by uncertainty in material properties as well as how predictions may be more properly characterized in a probabilistic sense. All of this helps to support condensed phase burning model development and engineering decision making. Before discussing calibration results for several materials, the thermal degradation model and the Bayesian inference tools are presented in the next section.

2 Theory

Bayesian calibration requires (1) a model, (2) data, and (3) an algorithm for evaluating Bayes' theorem. In this section, the TGA model to be calibrated and the Markov Chain Monte Carlo (MCMC) algorithm used are briefly described. The TGA data used for calibration will be discussed in Section 3.

2.1 Reaction Models

Burning models typically assume a serial reaction mechanism for the thermal degradation reactions. In a TGA experiment, it is generally assumed that the sample is small enough such that temperature and concentration gradients are approximately zero. If this assumption is valid, then the state of the system may be described by temperature, T , and composition. Typically, the composition of combustible materials is quantified in terms of the mass densities of some number, N , of notional components. These components do not typically represent distinct chemical species, but rather some composite material whose reactivities and thermophysical properties may be readily measured. Reaction rates are commonly assumed to be directly proportional to the bulk density of the reacting component. Since TGA is assumed to be spatially lumped, the control volume may be held constant such that the composition is described in terms of the masses of the N components present in the system, m_i . When the temperature is prescribed—as in TGA, the dynamics of the system are governed by

$$\frac{dm_1}{dt} = -k_1 m_1 \quad (1)$$

$$\frac{dm_i}{dt} = -k_i m_i + v_{i-1} k_{i-1} m_{i-1}, \quad i = 2, \dots, N-1 \quad (2)$$

$$\frac{dm_N}{dt} = v_{N-1} k_{N-1} m_{N-1} \quad (3)$$

where k_i and v_i are the temperature dependent rate constant and the residual mass fraction corresponding to reaction i , respectively. In this model, the initial conditions for the system of ordinary differential equations (ODEs) are $m_1(t=0) = m_0$ and $m_i(t=0) = 0$, $i > 1$ where m_0 is the initial mass in the TGA sample pan. The rate constants will be assumed

to have an Arrhenius form such that $k_i = A_i \exp(-E_i/RT)$. In a TGA experiment, only the total mass, $m = \sum_i m_i$ is measured. In this paper, mass is non-dimensionalized by the initial sample mass as $w \equiv m/m_0$.

Equations 1-3 describe a family of models depending on the number of components. Each member of this model family may be labeled by the number of reactions, $N_r \equiv N - 1$. Each reaction is specified by three parameters: A_i , E_i , and v_i . There are thus a total of $3N_r$ parameters in the model. A constraint on the residual mass fractions is imposed by the final non-dimensional mass as $w_f = \prod_i v_i$. Therefore, the total number of parameters to be calibrated for model N_r is $3N_r - 1$. Note that in the following, the pre-exponential will be represented as $\log A_i$ for convenience.

2.2 Bayesian Inference

Bayesian inference is a process whereby a probabilistic characterization of knowledge is updated rationally based on experimental data. In this paper, all of the calibration parameters, θ , are continuously varying, and so the appropriate probabilistic characterization is a PDF, $p(\theta)$. The foundational principle in Bayesian inference is Bayes' theorem. Bayes' theorem is a straightforward consequence of the product rule of probability theory, and relates $p(\theta)$ (referred to as the prior) to the belief about θ given some set of experimental data, d , called the posterior $p(\theta|d)$. Using this notation, Bayes' theorem may be stated as

$$p(\theta|d) \propto p(d|\theta) \times p(\theta) \quad (4)$$

where the constant of proportionality is $1/p(d)$ and is excluded for simplicity. The likelihood function, $p(d|\theta)$, is just the probability of observing the data d given that the model is true and its parameters equal θ .

Computation of the posterior first requires specification of the prior PDF and the likelihood function. In most applications, there is considerable uncertainty in θ . If there is no reason to prefer one range of values of θ to another, then the prior PDF should be represented by a uniform distribution $p(\theta) = (\theta_u - \theta_l)^{-1}$ for $\theta_l < \theta < \theta_u$ and $p(\theta) = 0$ elsewhere. For computational purposes, bounds on the prior must be specified. Sometimes this is possible by hard physical constraints (e.g., positive mass). Other times, judgement is required based on some degree of experience with the uncertain parameters. This reliance on judgement is not as problematic as it might sound since the appropriateness of the selected bounds may be evaluated after the posterior has been computed. Specifically, if the posterior appears to be significantly truncated at the bounds of the assumed uniform prior, then it should be reevaluated with broadened bounds.

The choice of the likelihood function depends on what the expected data should look like if the parameters assume the value θ . The data for TGA is a set of residual mass fractions at a corresponding set of temperatures, $d = \{(T, w)_k\}$. Assuming the model is true, then it is reasonable to expect a Gaussian error in the measurements centered on the model predicted value. That is, if $w^*(T|\theta)$ is the solution of the model assuming θ , then the likelihood function is

$$p((T, w)_k|\theta) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(w_k - w_k^*)^2}{2\sigma^2}\right], \quad \forall k \quad (5)$$

where $w_k^* \equiv w^*(T_k)$ and σ is the unknown standard deviation. If the data are independent, then the likelihood function for the entire data set is just the product of Eq. 5 over all data

points, k :

$$P(\{(T, w)_k\} | \theta) = \prod_k P((T, w)_k | \theta) \quad (6)$$

In using Eq. 6 as the likelihood function, it is implicitly assumed that the model predictions, w_k^* , are the actual values of the residual mass in the pan. The likelihood function returns the probability of measuring the actually observed data under the assumption that the measured data are normally distributed about the actual instantaneous mass in the pan.

In a limited sense, computing the posterior PDF is straightforward given the data and the choice of an appropriate likelihood function. However, as the dimensionality of the problem increases, straightforward Monte Carlo simulations become infeasible. Fortunately, the development of Markov Chain Monte Carlo (MCMC) algorithms has made the application of Bayesian inference feasible to many engineering problems [19]. An MCMC algorithm generates a Markov chain whose equilibrium distribution is the posterior. Thus, the values of θ generated by MCMC are samples from the desired posterior, $p(\theta|d)$. A Markov chain is a random process in which the steps in the process depend only on the immediately previous step. In the context of MCMC, the relevant Markov chain is a sequence of values of θ . The most common algorithm for performing MCMC is the Metropolis-Hastings (MH) algorithm [20, 21]. In the MH algorithm a set of candidate values for θ_{i+1} are sampled from some proposal distribution about θ_i . The candidate parameter values are then accepted with some probability defined by what is called the acceptance distribution. In Bayesian calibration, the acceptance distribution equals one (the candidate is always accepted) if the posterior value for θ_{i+1} is greater than that for θ_i . If the candidate is rejected, then $\theta_{i+1} = \theta_i$. Appropriate selection of the proposal and acceptance distributions assures that the generated Markov chain asymptotically approaches the equilibrium distribution, which is the posterior for Bayesian calibration. In this paper, MCMC is accomplished using an adaptive MH algorithm [22] with PyMC [23], which is a Python module for Bayesian inference.

3 Calibration

Kinetic parameters were calibrated using MCMC for three materials: high-impact polystyrene (HIPS), bisphenol-A polycarbonate (PC), and poly(vinyl chloride) (PVC). The TGA data for HIPS is from Stoliarov et al. [24]. The TGA data for the two charring polymers, PC and PVC, is also from Stoliarov et al. in [25]. For all three materials, calibrations were performed using data obtained at heating rates, β , of 3 K/min and 10 K/min. Note that in these references, TGA experiments were also performed at a heating rate of 30 K/min. However, at higher heating rates, the assumption that the sample is isothermal becomes questionable. It is generally recommended that for samples on the order of 10 mg that heating rates be kept less than or equal to 10 K/min [26]. All of these experiments were performed in nitrogen to prevent oxidation reactions. The calibration data set consists of three replicate experiments for each material/heat rate combination.

The lower and upper bounds for the kinetic parameters in Eqs. 1-3 that were used to construct the priors in this paper are listed in Table 1. For the two Arrhenius parameters, $\log A_i$ and E_i , the bounds were chosen by trial and error. That is, it was found that the posterior PDFs were all approximately zero outside of these bounds. The lower and upper bounds on the residual mass fractions v_i are simply logical bounds. Reactions cannot destroy mass, so $v_i \geq 0$, and reactions cannot create mass, so $v_i \leq 1$.

Simulation of TGA data, $w^*(T)$, were performed by using explicit Euler time integration of Eqs. 1-3 with a time step on the order of 0.01 s. It was found that such an algorithm

Table 1 Uniform prior bounds.

Parameter	Lower Bound	Upper Bound
$\log(A_i, 1/s)$	1	70
$E_i, \text{kJ/mol}$	40	800
ν_i	0	1

was robust and accurate within the range of model parameters studied. Since details of the TGA trace at large times are not important, negative masses were prevented by brute force. Specifically, if the model predicts negative values of the mass of a component in a given time step, the change in mass in that time step was automatically corrected so that the mass equaled zero at the end of the time step. The other rate terms in the ODEs were corrected in a consistent manner.

The number of MCMC samples was chosen by trial and error until the sample trace appeared well-mixed and the autocorrelation between samples was small. For a single reaction step model, $N_r = 1$, the number of TGA simulations required was relatively small. For three reactions, $N_r = 3$, the number of simulations was much larger and generation of the complete Markov Chain took up to one week to complete on a 3 GHz, 8 Gb workstation. All computations were performed in serial. For the $N_r = 3$ cases, a total of 400 million simulations were performed, the first 200 million were discarded—to eliminate any observable drift in the Markov chain traces, and every 10 thousandth simulation thereafter was preserved in construction of the posterior—to destroy any residual correlation between samples as observed in plots of sample autocorrelation. This resulted in a total of 20 000 samples from the posterior. These samples were used to construct a PDF by kernel density estimation (KDE) using Scott’s rule [27] to determine the bandwidth.

3.1 High-impact Polystyrene

In TGA experiments, HIPS produces a small amount of residue. Since this residue is non-carbonaceous, HIPS is generally treated as a non charring material—i.e., a burning slab of HIPS will not produce a protective char layer. As shown in Fig. 1, the thermal degradation of HIPS closely follows a single Arrhenius step. Accordingly, Stolarov et al. [24] assumed a single reaction and determined the Arrhenius parameters by a linear fit of the natural log of a calculated rate constant to the inverse of the temperature. The linear fit used TGA data from heating rates of 3 K/min, 10 K/min, and 30 K/min, and the resultant kinetic parameters tabulated in Table 2. The quality of this fit can be visualized by comparing the ‘x’ (model) and ‘o’ (experiment) markers in Fig. 1.

Using the HIPS data as input, kinetic parameters were calibrated for serial reaction mechanisms with one ($N_r = 1$) and two ($N_r = 2$) reaction models. The results of an MCMC calibration may be visualized in many ways. The TGA model prediction using the calibrated kinetic parameters as inputs can be compared to the experimental data used in MCMC simulation. Since the result of an MCMC simulation is a set of samples from the posterior PDF, the calibrated kinetic parameters should be thought of in probabilistic terms. Hence, the exact values of the kinetic parameters are uncertain. Propagation of this uncertainty through models will be further discussed in Section 4, but the simplest way to predict model results based on probabilistically characterized parameters is to perform the simulation using the mode of the posterior PDF. The mode is a reasonable point estimate of the parameters’ values since it is the most probable set of property values. The mode of the posterior was

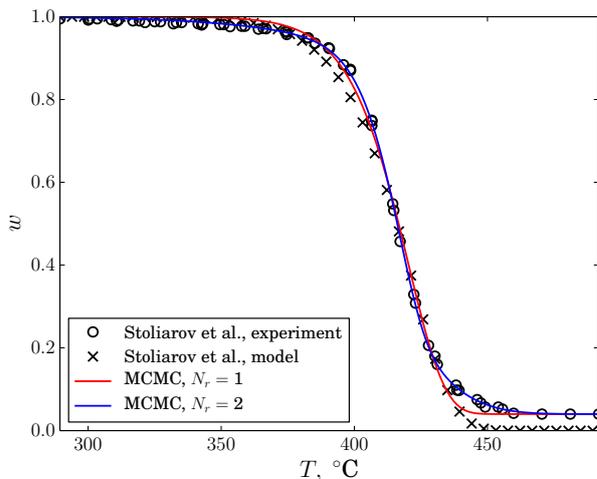


Fig. 1 Fits to TGA data for HIPS obtained at a heating rate of 3 K/min.

found by evaluating the KDE PDF at all of the MCMC posterior samples and finding the maximum value. Simulated TGA of HIPS at a heating rate of 3 K/min based on the posterior modes from the calibrated kinetic parameters using single and two-step reaction models are plotted in Fig. 1. As would be expected, the two-step reaction model provides a better fit to the data but the improvement gained by the additional reaction appears to be marginal. The posterior mode assuming a single-step provides a better fit to the data than the calibrated single-step kinetics of [24]. The better fit at high temperatures is simply because the model of the present paper takes into account the residual mass, but the better fit at temperatures of 370 °C to 400 °C could be attributed to either MCMC being a better fitting procedure or the fact that additional heating rate data is being considered in [24].

A second way to examine the MCMC results is to visualize the posterior PDFs. The univariate posteriors of the activation energy assuming a single reaction step for heating rates of 3 K/min and 10 K/min are plotted in Fig. 2. Also shown are the modes determined from the KDE of the bivariate PDF. Both distributions are approximately symmetric, and there is significant overlap indicating some degree of consistency between the two data sets with respect to the model. The point estimate of [24] is well into the lower tails of both distributions. This indicates that a good visual fit to the data can be achieved with relatively improbable kinetic parameters.

Since the one-step reaction model involves only two kinetic parameters, the bivariate distribution is the total PDF. Bivariate posteriors are plotted as contour plots in Fig. 3 for both heating rates. Note that the contours are not distinguishable since the distribution is so narrow. Also shown in Fig. 3 are the modes of these distributions and the point estimate of [24]. The distributions are extremely narrow and show a strong correlation between the pre-exponential and the activation energy. This correlation is often referred to as “kinetic compensation” [28, 9]. It is noteworthy that, while this compensation effect is strong, distinct peaks do exist. It is therefore not justifiable to assume any point along the line of correlation. However, even though distinct peaks exist in the PDF, it is apparently not the case that all fitting techniques will find them as is evidenced by the point estimate of [24]. In fact, because the posterior PDF is so narrow, it is likely that most local search algorithms will not

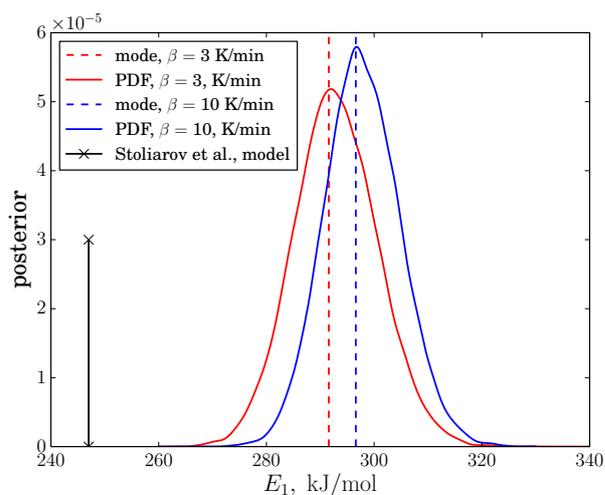


Fig. 2 Univariate posterior PDFs for HIPS assuming one reaction.

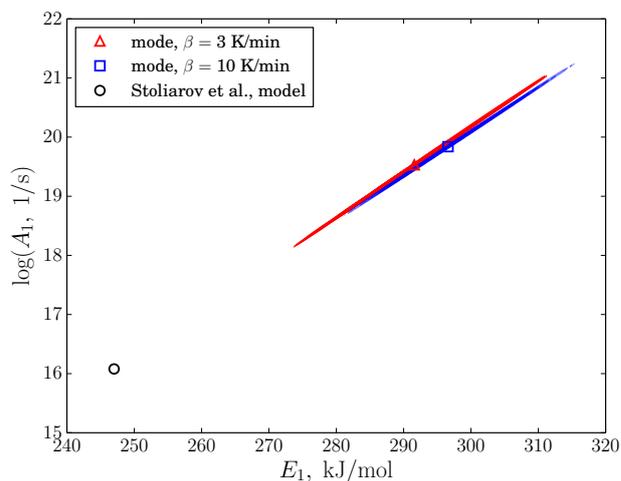


Fig. 3 Bivariate posterior PDFs for HIPS assuming one reaction.

in fact find the true posterior mode—that is, the true “best fit” to the TGA data. Therefore, it is advisable to perform a more exhaustive global fit procedure, such as that presented in this paper. While the distributions implied by the 3 K/min and 10 K/min heating rate TGA data are close, they do not perfectly overlap. Such differences indicate that the lumped, serial reaction model described in Eqs. 1-3 is not perfectly describing the physics of the TGA experiment. Some degree of imperfection is acceptable in engineering applications, and the degree to which approximation is allowed should be determined by the engineering application.

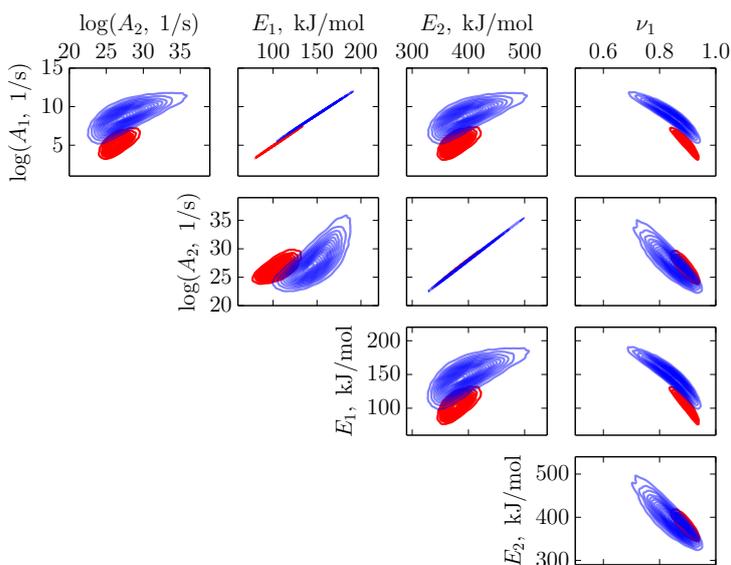


Fig. 4 Bivariate posterior PDFs for HIPS assuming two reactions. Red and blue contours correspond to 3 K/min and 10 K/min TGA data, respectively.

Although a single reaction model seems to provide a reasonable fit to the TGA data, polymer decomposition involves many more reactions. It is worthwhile therefore to consider higher-order models. To this end, a two reaction model, $N_r = 2$, was considered. The TGA simulation based on the mode of the two reaction posterior is plotted in Fig. 1, and it is apparent that the additional reaction results in a better fit at higher temperatures. The bivariate posteriors for the two reaction model are plotted in Fig. 4. Once again, significant correlation is observed between the activation energies and pre-exponentials within both reactions. While correlations are observed amongst the other parameters, the inter-reaction posteriors (e.g., E_1 versus E_2) are much broader. Note also that the correlations between reactions 1 and 2 are positive. In particular, consider the top left plot of $\log A_1$ versus $\log A_2$. As the rate of reaction 1 increases so does the rate of reaction 2. This makes sense because if the rate of reaction 1 increased without a corresponding rate of increase in reaction 2, a plateau would eventually develop in the TGA trace, $w(T)$, at the point where component 1 is exhausted but component 2 has not yet begun to degrade. Alternatively, if the rate of reaction 2 increased without a corresponding increase in the rate of reaction 1, a sharper drop-off in the TGA trace would develop as w approaches its final value. Similar trade-offs result in the correlations between the other inter-reaction kinetic parameters. The residual mass fraction from reaction 1 shows a significant negative correlation with the four other kinetic parameters. The source of this correlation is not clear.

3.2 Bisphenol-A Polycarbonate

Unlike HIPS, PC forms a significant char layer upon burning. This char formation is manifested in TGA experiments as a final residual mass fraction of around 0.21. A plot of 3

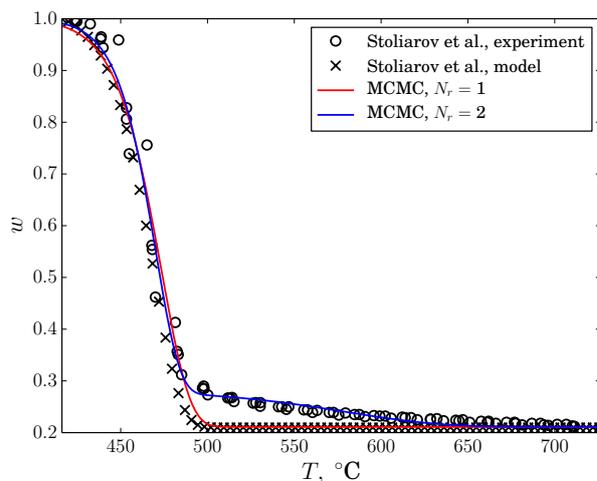


Fig. 5 Fits to TGA data for PC obtained at a heating rate of 3 K/min.

K/min TGA data for PC from [25] is provided as the ‘o’ markers in Fig. 5. This TGA trace has a significantly different shape from that of HIPS. Although the mass is initially lost with an Arrhenius type shape, at large temperatures, there is a long tail representing a gradual mass loss. A two reaction model could account for this, but since the mass loss associated with temperatures above 500 °C is relatively small, Stoliarov et al. [25] choose to model the pyrolysis of PC with a single reaction step. Stoliarov et al. [25] determined the Arrhenius parameters of this single-step model by fitting to the 3 K/min TGA data, and the reported values of these parameters are listed in Table 3.

Arrhenius parameters were calibrated using MCMC assuming both one-step and two-step models. The modes of the resultant posteriors were used to simulate the 3 K/min TGA trace and the results are plotted in Fig 5. The single-step model results in a fit similar to that of [25]. Naturally, a better fit is obtained with two reactions, and it is seen that the long tail observed in the TGA data is fairly accurately captured by the two reaction model.

The bivariate posteriors assuming a single reaction and calibrated against 3 and 10 K/min TGA data are plotted in Fig. 6. Unlike the case for HIPS, the literature point estimate from [25] falls within the more probable region of the 3 K/min posterior. Also, there is significant divergence between the posteriors obtained from the two heating rates. This is further indication that a single-step model is not appropriate for PC.

The two-step model posteriors are plotted in Fig. 7. In contrast to the two reaction bivariate plots for HIPS, there is relatively close agreement between the posteriors obtained at the two heating rates. This is evidence that an $N_r = 2$ model is consistent between the two data sets, a moderate form of validation. Kinetic compensation is observed for both reactions, but reaction 2 has a broader positive correlation between A_2 and E_2 . Overall, the posteriors for reaction 2 parameters are much broader than those for reaction 1. This is due to the fact that the residual mass from reaction 1 is small—i.e., $v_1 \approx 0.3$. Consequently, model predictions are less sensitive to reaction 2 kinetic parameters since reaction 2 accounts for a relatively small proportion of the total mass loss. In comparison with HIPS, the reaction 1 residual mass fraction does not show a strong correlation with the kinetic parameters. Again, this is because the TGA data indicate a relatively small value and so the exact shape of the

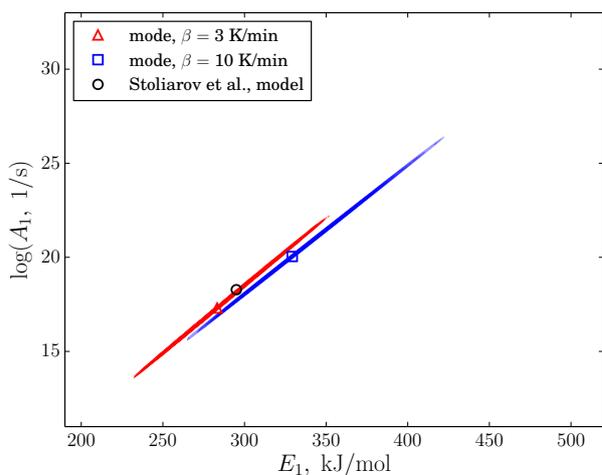


Fig. 6 Bivariate posterior PDFs for PC assuming one reaction.

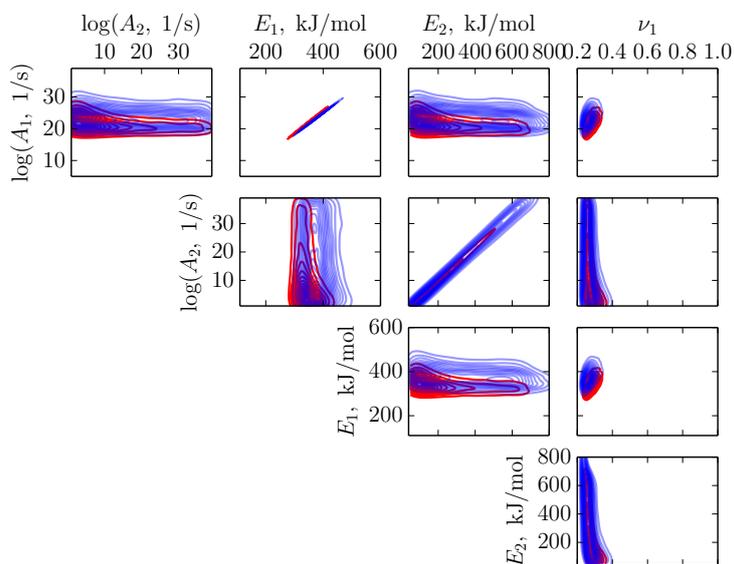


Fig. 7 Bivariate posterior PDFs for PC assuming two reactions. Red and blue contours correspond to 3 K/min and 10 K/min TGA data, respectively.

curve after the first reaction is relatively unimportant as long as the transition point between the two reactions is maintained. Consequently, there is even less correlation between the residual mass fraction and the reaction 2 kinetic parameters.

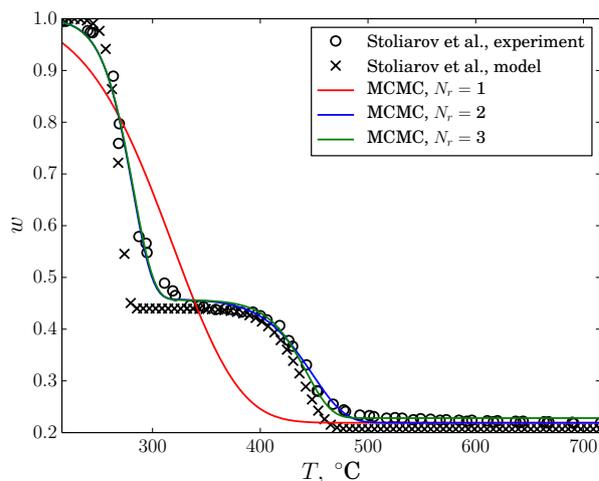


Fig. 8 Fits to TGA data for PVC obtained at a heating rate of 3 K/min.

3.3 Poly(Vinyl Chloride)

The final material considered in this paper is PVC. Not only does PVC form a significant char layer, but the TGA data indicate two distinct mass loss peaks with a possible third reaction corresponding to small mass loss at temperatures greater than 500 °C. These characteristics of PVC mass loss are shown in Fig. 8. Stoliarov et al. [25] modeled the thermal degradation of PVC with two serial reactions, and the simulation of this model is represented by the ‘x’ markers in Fig. 8. A summary of the calibrated kinetic parameters is provided in Table 4.

The quality of the MCMC calibrations assuming one, two, and three reaction steps may be assessed by plotting the TGA traces predicted using the posterior mode parameters. These traces are shown in Fig. 8. The inadequacy of a single reaction step is apparent. However, the two and three reaction step models are nearly indistinguishable. It is interesting to note that the MCMC fits in this case are significantly better than the fits obtained in [25] using a more ad hoc approach to calibration. This is in contrast to the results for HIPS and PC where the advantages of using the MCMC approach were less apparent.

Bivariate posteriors assuming a two reaction model are plotted in Fig. 9. The point estimates of [25] are shown as circles. As with PC, the two-reaction model results in close agreement between the posteriors derived using both the 3 K/min and 10 K/min TGA data. In contrast to PC, the breadth of the distributions in A_2 and E_2 is much narrower indicating that the data better inform these two parameters. Kinetic compensation is observed for both reactions, and the point estimates of [25] fall along the kinetic compensation lines even though the estimates for the first reaction are well outside of the most probable region. The correlations between the two reactions (e.g., A_1 vs. E_2) are relatively small. This is due to the fact that the two reactions are well-distinguished, the first reaction finishing well before the second reaction begins. Similarly, the residual mass fraction, v_1 , shows little correlation with any of the Arrhenius parameters. This parameter essentially fixes the vertical location of the plateau in w between the two reactions and is therefore relatively insensitive to the exact shape of these two portions of the mass trace.

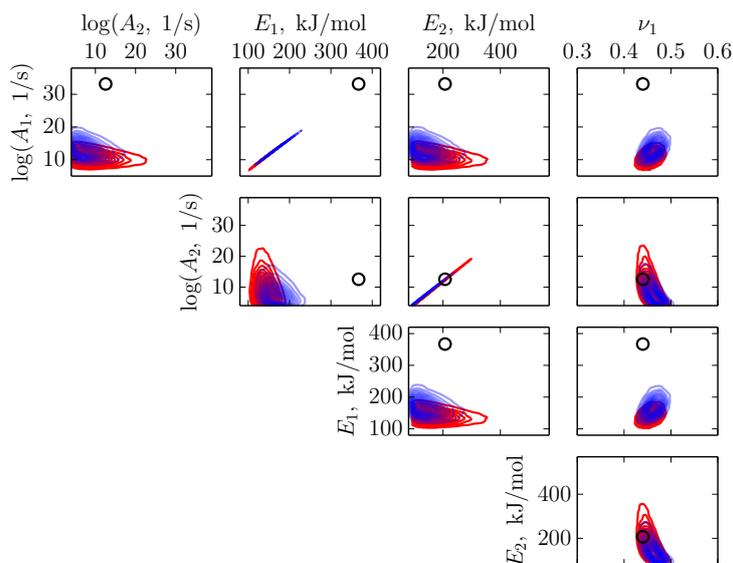


Fig. 9 Bivariate posterior PDFs for PVC assuming two reactions. Red and blue contours correspond to 3 and 10 K/min TGA data, respectively. Circles correspond to estimates of Stoliarov et al. [25].

Finally, the bivariate posteriors resulting from the MCMC calibration assuming a three reaction model are plotted in Fig. 10. These posteriors are considerably more complicated than any of the previous results. This is attributable to the complexity of the model and the various possibilities for compensation effects between the 8 parameters. Despite this increased complexity, several patterns are apparent. First, as with all other cases, kinetic compensation between the intra-reaction Arrhenius parameters is observed. The breadth of these intra-reaction posteriors increase with their sequence in the serial reaction order. Second, there is much agreement between the posteriors obtained at the two heating rates even though this agreement is less pronounced as compared to the two reaction case. Third, the inter-reaction posteriors typically assume an ‘L’-shape. For example, consider the plot of E_1 vs. E_3 . The most probable combination of these two parameters is the case of both being slightly less than 200 kJ/mol implying that both reactions begin at about the same temperature. If either E_1 or E_3 is kept constant at this value, then the other may likely assume a fairly broad range of larger values. That is, as long as one of the two reactions occurs at a fairly low temperature, the other reaction is relatively insensitive to temperature. This seems unreasonable unless the residual mass fractions are taken into account. For instance, consider the case where E_1 is kept around 150 kJ/mol while E_3 is increased. In this case, reaction 1 is happening relatively fast while reaction 3 is becoming increasingly slow. Now consider the posteriors of E_3 versus ν_1 and ν_2 . In both cases, as E_3 increases the more probable residual mass fractions become smaller implying that the actual mass loss dependent on reaction 1 increases. Consequently, in the case where reaction 1 is fast, reaction 3 may be relatively slow as long as the majority of the actual mass lost is associated with reaction 1. A similar analysis reveals that reaction 1 may be slow relative to reaction 3 as long as reactions 2 and 3 account for the majority of the actual mass lost. Figure 10 also reveals the

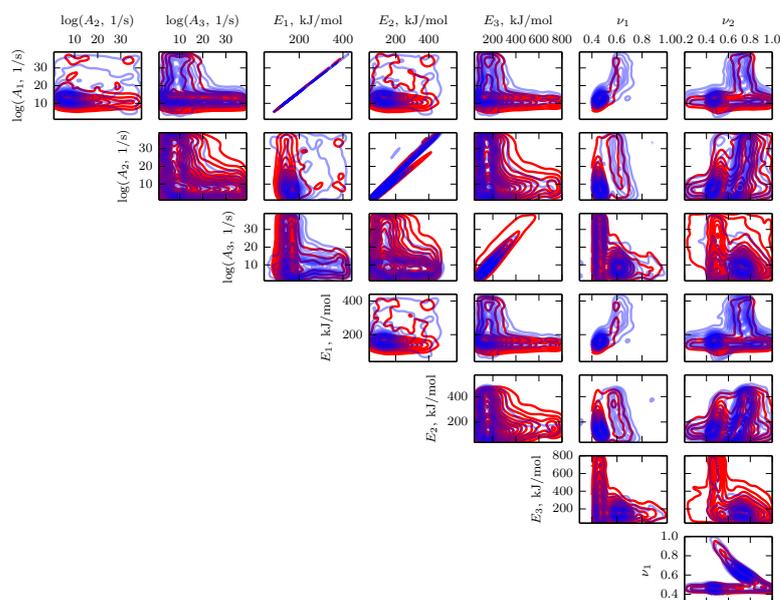


Fig. 10 Bivariate posterior PDFs for PVC assuming three reactions. Red and blue contours correspond to 3 and 10 K/min TGA data, respectively.

correlation between the residual masses generated by reactions 1 and 2. This correlation is visualized in the posterior in the lower right corner panel. Note that ν_1 has a definite lower bound of around 0.4. This value is fixed by the plateau in w visible in Fig. 8. In the cases where $\nu_1 \approx 0.4$, reaction 1 is complete before the other two reactions begin. The relative proportions at which mass is lost between reactions 2 and 3 is not too significant in such cases since the first reaction accounts for the majority of the mass loss. Therefore, as is seen in Fig. 10, a broad range of values for ν_2 are probable if $\nu_1 \approx 0.4$. Notice however that in cases where $\nu_1 > 0.5$, there is a strong negative correlation between ν_1 and ν_2 . In other words, if less mass is lost in reaction 1 (ν_1 increasing) then more mass must be lost in reaction 2 (ν_2 decreasing) in order to account for the first mass loss peak—i.e., while $w > 0.4$. In these cases, reaction 1 and reaction 2 correspond to the first mass loss peak while reaction 3 corresponds to the second (at $T \approx 450$ °C). Although these results are interesting, and shed some light on the types of tradeoffs that occur in serial reaction networks, it is apparent from the TGA plot, Fig. 8, that the two reaction model provides a sufficient fit to the TGA data and that the calibrated parameters are relatively consistent as a function of heating rate (Fig. 9). For these reasons, a two reaction model seems preferable to a three reaction model.

Table 2 Calibrated kinetic parameter posterior modes for HIPS.

	MCMC, 3 K/min	MCMC, 10 K/min	Stoliarov et al. [24]
$\log(A_1, 1/s)$	19.5	19.8	16.1
$E_1, \text{kJ/mol}$	292	297	247
$\log(A_1, 1/s)$	5.24	8.60	-
$\log(A_2, 1/s)$	26.4	27.2	-
$E_1, \text{kJ/mol}$	104	145	-
$E_2, \text{kJ/mol}$	378	390	-
ν_1	0.897	0.848	-

Table 3 Calibrated kinetic parameter posterior modes for PC.

	MCMC, 3 K/min	MCMC, 10 K/min	Stoliarov et al. [25]
$\log(A_1, 1/s)$	17.3	20.0	18.3
$E_1, \text{kJ/mol}$	283	330	295
$\log(A_1, 1/s)$	20.9	23.0	-
$\log(A_2, 1/s)$	3.83	6.99	-
$E_1, \text{kJ/mol}$	332	372	-
$E_2, \text{kJ/mol}$	112	157	-
ν_1	0.277	0.274	-

Table 4 Calibrated kinetic parameter posterior modes for PVC.

	MCMC, 3 K/min	MCMC, 10 K/min	Stoliarov et al. [25]
$\log(A_1, 1/s)$	10.6	12.6	33.1
$\log(A_2, 1/s)$	7.44	6.70	12.5
$E_1, \text{kJ/mol}$	140	161	367
$E_2, \text{kJ/mol}$	141	127	207
ν_1	0.458	0.466	0.44
$\log(A_1, 1/s)$	10.3	23.1	-
$\log(A_2, 1/s)$	10.8	12.7	-
$\log(A_3, 1/s)$	14.5	10.7	-
$E_1, \text{kJ/mol}$	137	269	-
$E_2, \text{kJ/mol}$	184	168	-
$E_3, \text{kJ/mol}$	715	182	-
ν_1	0.457	0.626	-
ν_2	0.500	0.725	-

3.4 Summary

In this section, Bayesian inference via MCMC was used to calibrate the parameters of one, two, and three step serial reaction networks for modeling the thermal degradation of three materials. The resultant posterior PDFs are useful for characterizing uncertainty, identifying parameter compensation effects, and possibly determining the appropriate level of model complexity. Although the results of this exercise are most faithfully characterized by the posterior PDFs plotted and discussed in the previous paragraphs, it is helpful to summarize the results in terms of point estimates for the various parameters. As mentioned previously, posterior modes are reasonable point estimates as they represent the most probable single point in the inferred parameter space. The MCMC derived posterior modes are tabulated in Tables 2-4 along with the point estimates of Stoliarov et al. [24,25]. In the next section, the parameter posterior distributions calibrated above are used to make predictions.

4 Validation and Prediction

Calibrated material properties serve as inputs to models. Ultimately, these models are to be used to make predictions to aid in engineering decision making. For example, a solid burning model could be used to assist simulations of fire growth in a building design and therefore help engineers decide on an appropriate level of fire protection. Before reliable predictions can be made, however, the model should be validated against experimental data as similar to the design scenario as possible. Both prediction and validation require the propagation of calibrated material properties through the relevant model.

In a typical deterministic calibration, point estimates for material property values allow for straightforward predictions. When these predictions do not agree with experimental observations it is not always clear whether the property estimates are wrong or if there is a significant modeling error. One advantage of Bayesian inference is that information regarding the uncertainty in the property estimate is retained. If this uncertainty is propagated during model validation, then model and parameter uncertainty may be distinguished. If this uncertainty is propagated during prediction, the decision maker has more information regarding the probabilities of various outcomes.

In this section, samples from the kinetic parameter posterior distributions presented in Section 3 are propagated through TGA and burning models. In all cases, the samples are simply taken from the Markov Chain.

4.1 Thermogravimetric Analysis

Although the MCMC simulations were calibrated against TGA data, it is not generally well understood how sensitive calibrated kinetic parameters are to heating rates. As was seen in Section 3, for each material and model combination, the posterior was slightly different for the two heating rates. Sometimes this difference was large, as in the case for HIPS with $N_r = 2$. In other cases, the posteriors resultant from the 3 K/min and 10 K/min heating rate TGA data were in significant agreement, as in the case for PVC with $N_r = 2$. Even though the posteriors differ, it is still possible that the TGA simulations based on these posteriors would be in significant agreement at both heating rates. That is, the TGA model predictions at different heating rates might be relatively insensitive to the parameter values within the range of the calibrated posteriors.

To better understand this, the posterior distributions obtained using 3 K/min TGA data were sampled and propagated through the TGA model at a heating rate of 10 K/min. For each combination of material, heating rate, and number of reactions (N_r), the first 100 samples of the Markov chain were propagated. Since the model is fast, many more samples could have been used, but it was found that 100 gave a reasonable coverage of the range of predicted TGA traces.

A plot of the 10 K/min TGA data for HIPS is provided in Fig. 11. The point estimate parameters from Stoliarov et al. [24] provide a reasonable fit to the data. Recall that these literature parameters were partially obtained using this 10 K/min data so this is not a case of pure prediction, but this simulation result was included for completeness and comparison with the actual predictions based on MCMC calibrated kinetic parameters. Also plotted in Fig. 11 are the predicted traces using posterior modes obtained from calibration to 3 K/min data assuming both one and two step reactions. It is apparent that the one reaction model provides a better fit to the data. This is an example of overfitting, and demonstrates the importance of not using more complex models than are necessary as this can result in poor

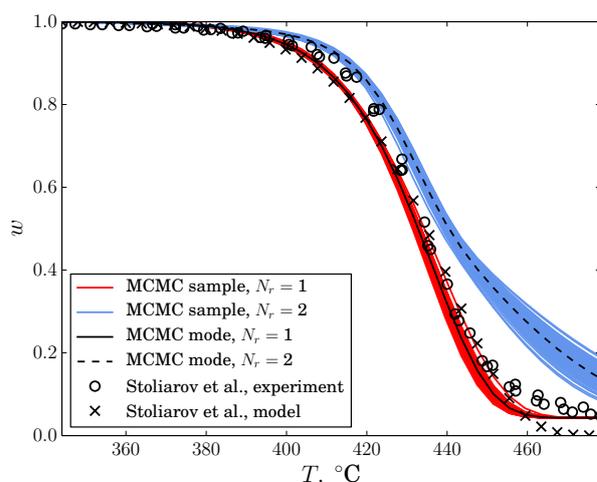


Fig. 11 Predictions of 10 K/min HIPS TGA data using parameters calibrated to 3 K/min TGA data.

predictions. The lines (there are many so that the lines appear as a shaded region) in Fig. 11 represent simulations generated using samples from the calibrated posteriors as input parameters. Note that the spread in the two reaction model is much greater than that in the one reaction model. This can partly be explained by the breadth in the two reaction posteriors shown in Fig. 4 as compared to that in the one reaction posterior shown in Fig. 3.

A similar plot is presented in Fig. 12 for PC. The one reaction MCMC mode results in a slightly better prediction than the point estimate of Stoliarov et al. [25]. This is in contrast to the case with HIPS where the point estimate of [24] performed slightly better. The difference in the case of PC is that both the MCMC mode and the point estimate of Stoliarov et al. are both true predictions, whereas in the case of HIPS only the MCMC mode was predictive. The two reaction MCMC mode give comparatively mixed results. Although, unlike the one reaction mode, it captures the $T > 500$ °C tail, it underpredicts the initial mass loss temperature by slightly more than the one reaction model. In this case, the spread between the two models is comparable. The predictive ability of these models, even in a relatively moderate extrapolation to a higher heating rate would appear to be weaker for PC than for HIPS. This might be attributed to the neglect of some relevant physics such as mass transfer.

The same scenario is considered for PVC, and the results are plotted in Fig. 13. For this case, two and three step serial reaction mechanisms are used. As with PC, the MCMC mode-based predictions performed better compared to the point estimate of Stoliarov et al. [25]. The three reaction mode captures the second mass loss dip ($T \approx 475$ °C) better than the two reaction mode. The spread in the sampled predictions covers the data for both models, but the breadth of w curves for the three reaction model is slightly larger. Because the data falls within the uncertainty of the predictions, it is likely that both models are appropriate. This is an example of how a model might be validated in a probabilistic sense.

The ability of MCMC calibrated kinetic parameters for predicting TGA has been investigated. It was found that point estimates of these parameters based on the posterior modes is more successful for predictions than more ad hoc point estimates typically presented in the literature. It was further demonstrated that a probabilistic characterization of the model pa-

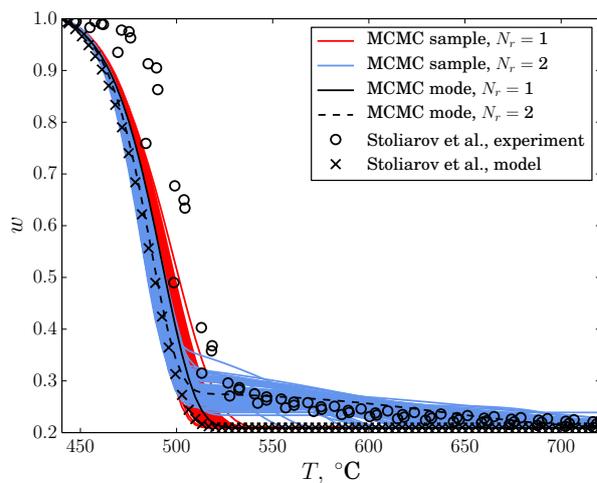


Fig. 12 Predictions of 10 K/min PC TGA data using parameters calibrated to 3 K/min TGA data.

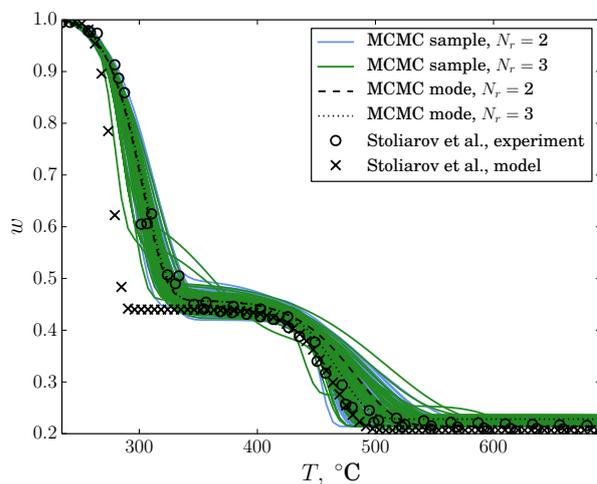


Fig. 13 Predictions of 10 K/min PVC TGA data using parameters calibrated to 3 K/min TGA data.

rameters gives a more detailed account of the predictive capability of the model by providing practitioners with the probabilities of various possible outcomes.

4.2 Burning Rate

Ultimately, TGA is an experiment used for property calibration. A more interesting engineering problem is the prediction of the burning rate of materials. Two relevant experiments for burning rate are gasification and cone calorimetry. In these experiments, heat and mass

transfer, in addition to thermal degradation chemistry, are significant contributing physical processes.

All simulations were carried out using the Fire Dynamics Simulator (FDS) version 6.1.0 [29], at SVN revision number 19823. Both gasification and cone calorimetry simulations relied on the FDS solid phase submodel, which allows for one-dimensional heat transfer, multicomponent thermophysical properties, and serial chemical reaction mechanisms. The properties for all cases were those suggested by Stoliarov et al. [24,25], and the FDS input files were based on those for the “FAA Polymers” case in the FDS Validation Guide [1]. Since property values are unknown for intermediate states, only the reaction models corresponding to the number of reactions in [24,25] were used. That is, for HIPS and PC only $N_r = 1$ was simulated, and for PVC only $N_r = 2$ was simulated. In future work, it would be valuable to see the effect of additional reactions on the FDS simulations using crude estimates of the other material properties for intermediate states. Although replicate burning rate experiments were performed, the data are not available, and so experimental uncertainty is not plotted in the following figures. Uncertainty in burning rate is significant, though, and can approach up to 25 % in the vicinity of the peak [25].

The first burning rate experiment considered was the gasification of HIPS using the data presented in [24] obtained using the NIST gasification apparatus [30]. The experimental data along with sampled predictions are plotted in Fig. 14 for a nominal external heating flux of 50 kW/m^2 . A total of 1000 samples were taken from each of the 3 K/min and 10 K/min posteriors and propagated through the FDS model. It is apparent that there is relatively little scatter in the gasification rate predictions based on the MCMC generated kinetic parameter posteriors. The MCMC posterior sampled kinetic parameters and the point estimate of Stoliarov et al. [24] all result in similar FDS simulation results. The 3 K/min samples predict a slightly faster gasification process than the 10 K/min samples, which in turn predict a slightly faster gasification than the point estimate of Stoliarov et al. All cases predict a peak gasification point that is slightly too low and too soon. The fact that the experimental data do not fall within the scatter implied by the uncertainty in the kinetic parameters indicates the presence of error in the model or in the other thermophysical properties input into the model.

The gasification simulations were relatively fast in terms of CPU time since no gas phase combustion was included. Therefore a large number of samples of kinetic parameters could be propagated through FDS. For the remaining two materials, predictions were made of cone calorimetry [31] in which it was necessary to include the effects of the above surface flame. Several scenarios for PC were studied in [25], but in the following only the 5.4 mm thick sample exposed to a heating rate of 50 kW/m^2 was considered. The plot of the measured and predicted heat release rate (HRR) is given in Fig. 15. Because the gas phase combustion slows down the simulation considerably, only 55 samples were taken from each of the two heating rate posteriors. As was the case with HIPS, the sampled solutions are tightly grouped and mostly agree with the result obtained using the point estimate of Stoliarov et al. The simulations based on the posterior calibrated against the 10 K/min TGA data tend to better predict the data for times greater than about 250 s relative to the predictions from parameters calibrated against the 3 K/min data. The MCMC sampled solutions all takeoff somewhat slower than the data and the literature point estimate, and they generally underpredict the first peak in HRR. Since the amount of energy contained in the material is the same for all solutions, this initial tendency to underprediction in HRR is made up for by a larger HRR at times above around 400 s. This observation is more in line with the experimental data as compared to the literature point estimate. Although all of the simulations agree reasonably well with the cone calorimetry data for small times, the divergence between model and

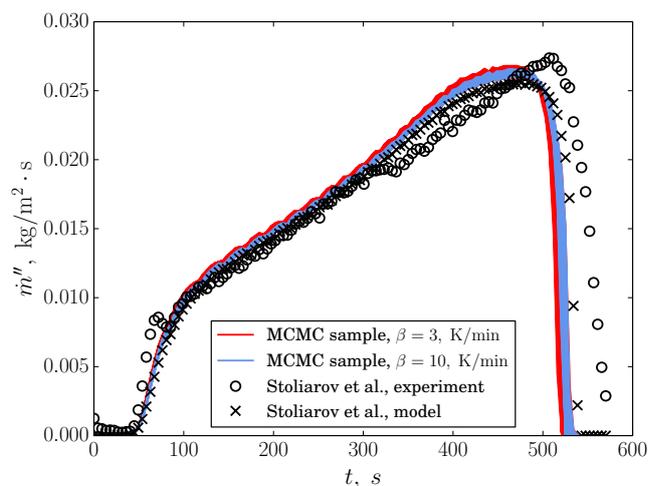


Fig. 14 Predictions of HIPS gasification experiments with an external heat flux of 50 kW/m^2 .

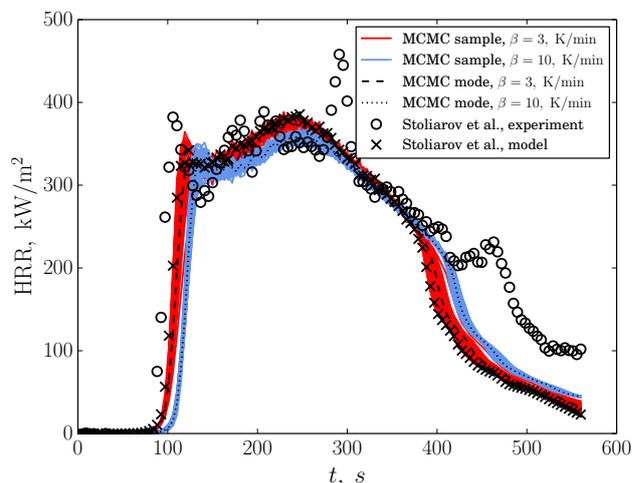


Fig. 15 Predictions of PC cone calorimetry experiments with an external heat flux of 50 kW/m^2 .

experiment at larger times must be attributed to model error or inaccuracies in the non-kinetic thermophysical properties used to describe PC in the FDS solid model. It is also observed that the peak HRR at about 300 s is not predicted by any of the simulations. This is possibly due to the fact that the model does not account for an apparent second reaction seen in the TGA data.

Finally, cone calorimetry of PVC was simulated for the case of a 6 mm thick sample exposed to a heat flux of 75 kW/m^2 [25]. The simulation times were long, and so only 65 samples from the MCMC generated kinetic parameter posteriors were propagated through FDS. The results, along with the experimental data, are plotted in Fig. 16. Similar to the

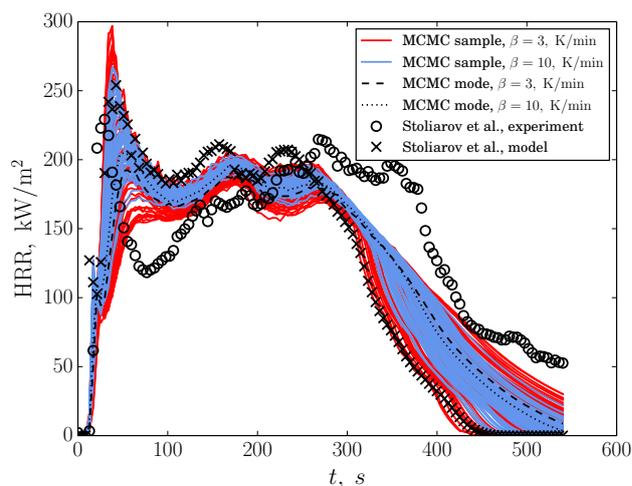


Fig. 16 Predictions of PVC cone calorimetry experiments with an external heat flux of 75 kW/m^2 .

TGA predictions for PVC, there is significantly more scatter in the predicted burning rate as compared to the other two materials. There is somewhat more scatter in the predictions based on the 3 K/min calibrated parameters. The MCMC sampled solutions generally provide better overall fits to the data as compared with the predictions made using the point estimate of Stoliarov et al. The point estimate solutions corresponding to the MCMC posteriors are seen to over predict the peak HRR time and under predict the peak HRR. At larger times, however, the MCMC posterior mode predictions are closer to the experimental HRR as compared to the point estimate solution of Stoliarov et al. [25].

It is also helpful to examine the scatter in FDS model predictions for specific quantities of interest. For fire safety applications, the peak HRR, HRR_{\max} , and the time at which this peak HRR occurs, t_{\max} are of particular interest. These quantities roughly characterize the severity of a fire fueled by the studied material. The values of these quantities predicted by the FDS model are plotted in Fig. 17 for each of the sets of kinetic parameters sampled from the posterior PDFs. The experimentally measured HRR_{\max} falls within the range of the sampled model predictions, but the experimentally measured t_{\max} is approximately 50 % less than the majority of the sampled predictions. It is also noteworthy that there exists a cluster of sampled FDS solutions with t_{\max} in the range of 160 to 190 s. This means that there are sets of kinetic parameters that provide reasonable fits to the TGA data that result in over predictions of the time to peak HRR by almost an order of magnitude. Such an over prediction could lead to engineering decisions that endanger property and human life. This is further demonstration of the importance of rigorous model validation in order to avoid making potentially dangerous engineering decisions.

5 Conclusions

In this paper, Bayesian inference methods were used to calibrate parameters for a simple class of thermal degradation models. The calibrations were performed using MCMC and TGA data at two heating rates. Resultant from these calibrations are samples from the poste-

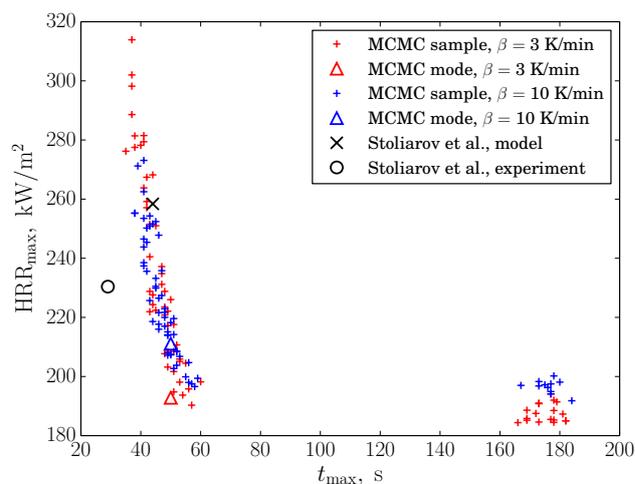


Fig. 17 Predicted maximum HRRs and times for PVC cone calorimetry.

rior distributions implied by the TGA data and the assumption of broad uniform prior PDFs. These samples were used to construct KDE estimates to the posterior PDFs. To demonstrate the value of this type of calibration, a subset of the MCMC generated posterior samples were propagated through TGA and FDS burning rate models.

The calibrations demonstrated that the proposed methodology is both feasible and valuable. It was shown to be feasible by the breadth of examples considered. Specifically, three materials were studied each at two heating rates each with two models for a total of 12 unique scenarios. The materials studied were sufficiently different to be representative of a broad range of realistic engineering polymers. The methodology was shown to be valuable by the increase in information provided by bivariate posterior PDFs. These PDFs reveal the uncertainty and parameter correlations implied by the TGA data. This information goes well beyond what is encoded in typical point estimate approaches to parameter calibration. The results indicate that overlap between posteriors obtained from different heating rate scenarios might be a measure of the adequacy of the model.

A further value of the Bayesian inference approach to parameter calibration is that it generates data that can be used to characterize uncertainty in model predictions. This uncertainty quantification allows for a more rigorous approach to model validation. In particular, if the uncertainty in the model prediction implied by the uncertainty in the kinetic parameter uncertainties covers the experimental data, then more accurate predictions could be made by more precisely measuring the kinetic parameters. This was the case, for instance, in the 10 K/min TGA predictions for PVC. On the other hand, if the sampled predictions are tightly grouped and do not cover the experimental data, then better future predictions will be more readily achieved by either improving the model to incorporate more physics or better characterizing other material properties used in the model. This was the case, for instance, in the gasification predictions for PC. In some cases, it was found that simpler calibrated models made better predictions than more complicated models. This is a warning against the dangers of overfitting. Finally, it was shown that point estimates for material properties in burning models can lead to faulty engineering decisions. This was the case for predictions of the burning rate of PVC in which it was found that some of the MCMC sampled

kinetic parameters, which provide reasonable fits to the TGA data, can result in an order of magnitude over prediction in the time to peak HRR.

The proposed methodology might be improved in several respects. First, better algorithms could be used for doing the MCMC. The Metropolis-Hastings algorithm was used in the present work, but more sophisticated algorithms have been developed that could open up the methodology to more complicated models. Second, other material properties relevant to condensed phase burning should be calibrated. For example, differential scanning calorimetry (DSC) data could be used to produce posterior PDFs for thermodynamic properties such as specific heat capacity and heats of reaction. This information, coupled with the kinetic information presented in the present paper, could be used to more thoroughly characterize the uncertainty in model predictions. Such a characterization would be extremely helpful for progress in modeling of condensed phase physics for fire problems as they give information on the relative amount of model error as opposed to material property measurement error. Third, alternative kinetic models could be explored. The present paper only calibrated serial reaction models. Some systems might be better characterized by an equal number of parallel reactions describing multiple competing components. Fourth, more detailed models of TGA and the burning rate experiments could be studied. In particular, heat and mass transfer could be modeled in the TGA to see if the difference in results at different heating rates are attributable to transport effects. Furthermore, oxidative reactions could be incorporated into both TGA and burning rate experiments. A final area for improvement in the methodology introduced in this paper would be to study the impact of calibrating against combined sets of TGA data—that is, in using TGA data from multiple heating rates to produce a single posterior. Such an approach might better characterize the kinetic parameter uncertainty and could possibly result in a different degree or uncertainty in model predictions.

References

1. K. McGrattan, S. Hostikka, R. McDermott, J. Floyd, C. Weinschenk, and K. Overholt. Fire dynamics simulator, technical reference guide volume 3: Validation. Technical report, NIST, 2014.
2. K. McGrattan, S. Hostikka, R. McDermott, J. Floyd, C. Weinschenk, and K. Overholt. Fire dynamics simulator, technical reference guide volume 1: Mathematical model. Technical report, NIST, 2014.
3. C. Lautenberger and C. Fernandez-Pello. Generalized pyrolysis model for combustible solids. *Fire Safety Journal*, 44:819–839, 2009.
4. Stanislav I. Stolarov and Richard E. Lyon. Thermo-kinetic model of burning. Technical Report DOT/FAA/AR-TN08/17, FAA, 2008.
5. C. Lautenberger, G. Rein, and C. Fernandez-Pello. The application of a genetic algorithm to estimate material properties for fire modeling from bench-scale fire test data. *Fire Safety Journal*, 41:204–214, 2006.
6. M. Chaos, M.M. Khan, N. Krishnamoorthy, J.L. de Ris, and S.B. Dorofeev. Evaluation of optimization schemes and determination of solid fuel properties for CFD fire models using bench-scale pyrolysis tests. *Proceedings of the Combustion Institute*, 33:2599–2606, 2011.
7. C. Lautenberger and C. Fernandez-Pello. Optimization algorithms for material pyrolysis property estimation. In *Fire Safety Science: 10th International Symposium*, College Park, MD, 2011.
8. G. Rein, C. Lautenberger, C. Fernandez-Pello, J.L. Torero, and D. Urban. Application of genetic algorithms and thermogravimetry to determine the kinetics of polyurethane foam in smoldering combustion. *Combustion and Flame*, 146:95–108, 2006.
9. M.C. Bruns, J.H. Koo, and O.A. Ezekoye. Population-based models of thermoplastic degradation: Using optimization to determine model parameters. *Polymer Degradation and Stability*, 94(6):1013 – 1022, 2009.
10. A. Gelman, J.B. Carlin, H.S. Stern, and D.B. Rubin. *Bayesian Data Analysis*. Chapman & Hall/CRC, Boca Raton, second edition, 2003.
11. D.S. Sivia. *Data Analysis*. Oxford University Press, Oxford, second edition, 2006.
12. R.D. Moser, G. Terejanu, T.A. Oliver, and C.S. Simmons. Validating the prediction of unobserved quantities. Technical Report 12-32, The Institute for Computational Engineering Sciences, 2012.

13. K.J. Overholt and O.A. Ezekoye. Quantitative testing of fire scenario hypotheses: A bayesian inference approach. *Fire Technology*, pages 1–33, 2014.
14. K.J. Overholt. *Forward and Inverse Modeling of Fire Physics Towards Fire Scene Reconstructions*. PhD thesis, The University of Texas at Austin, 2013.
15. S.I. Stoliarov, N. Safranov, and R.E. Lyon. The effect of variation in polymer properties on the rate of burning. *Fire and Materials*, 33:257–271, 2009.
16. M. Chaos. Application of sensitivity analyses to condensed-phase pyrolysis modeling. *Fire Safety Journal*, 61:254–264, 2013.
17. N. Bal and G. Rein. Relevant model complexity for non-charring polymer pyrolysis. *Fire Safety Journal*, 61:36–44, 2013.
18. G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6:461–464, 1978.
19. W.M. Bolstad. *Understanding Computational Bayesian Statistics*. Wiley, Hoboken, NJ, 2010.
20. N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, and A.H. Teller. Equation of state calculations for fast computing machines. *The Journal of Chemical Physics*, 21:1087–1092, 1953.
21. W.K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57:97–109, 1970.
22. H. Haario, E. Saksman, and J. Tamminen. An adaptive metropolis algorithm. *Bernoulli*, 7:223–242, 2001.
23. A. Patil, D. Huard, and C.J. Fonnesbeck. Pymc: Bayesian stochastic modelling in python. *Journal of Statistical Software*, 35:1–81, 2010.
24. S.I. Stoliarov, S. Crowley, R.E. Lyon, and G.T. Linteris. Prediction of the burning rates of non-charring polymers. *Combustion and Flame*, 156:1068–1083, 2009.
25. S.I. Stoliarov, S. Crowley, R.N. Walters, and R.E. Lyon. Prediction of the burning rates of charring polymers. *Combustion and Flame*, 157:2024–2034, 2010.
26. R.E. Lyon, N. Safranov, J. Senese, and S.I. Stoliarov. Thermokinetic model of sample response in nonisothermal analysis. *Thermochimica Acta*, 545:82–89, 2012.
27. D.W. Scott. *Multivariate Density Estimation: Theory, Practice, and Visualization*, pages 125–190. John Wiley & Sons, New York, 1992.
28. J. Ceamanos, J.F. Mastral, A. Millera, and M.E. Aldea. Kinetics of pyrolysis of high density polyethylene. comparison of isothermal and dynamic experiments. *Journal of Analytical and Applied Pyrolysis*, 65:93–110, 2002.
29. K. McGrattan, S. Hostikka, R. McDermott, J. Floyd, C. Weinschenk, and K. Overholt. Fire dynamics simulator, user’s guide. Technical report, NIST, 2014.
30. P.J. Austin, R.R. Buch, and T. Kashiwagi. Gasification of silicone fluids under external thermal radiation part i. gasification rate and global heat of gasification. *Fire and Materials*, 22:221–237, 1998.
31. American Society for Testing and Materials, West Conshohocken, Pennsylvania. *ASTM E 1354-04a, Standard Test Method for Heat and Visible Smoke Release Rates for Materials and Products Using an Oxygen Combustion Calorimeter*, 2007.