Efficient Simulation of a Reaction-Diffusion System with a Fast Reaction in the Asymptotic Limit

Aaron Churchill, Guan Wang, Matthias K. Gobbert, and Thomas I. Seidman
Department of Mathematics and Statistics, University of Maryland, Baltimore County,
1000 Hilltop Circle, Baltimore, MD 21250, U.S.A.

Abstract

We study a reaction-diffusion system of three chemical species, where two chemicals react with a much faster reaction rate than the other reaction in the model. We are interested in the asymptotic limit as the fast reaction rate becomes infinite. This forces the reaction interface to have an asymptotically small width with asymptotically large height. This interface is moving in time and causes interior layers that are progressively more challenging and costly for numerical simulations of the three species model, as the singularity becomes sharper with larger reaction rates. But in the asymptotic limit, an equivalent two component model can be defined that is significantly cheaper computationally and allows for effective studies for the model. The equivalence is demonstrated by the analytical definition of the two component model and by comparing numerical results to ones for the three species model with progressively larger reaction rates, which also demonstrate the computational efficiency. The state-of-the-art finite element package COMSOL Multiphysics is used for the simulations, thus also showing a practical way how to handle and visualize moving interior layers in reaction-diffusion systems. COMSOL is popular in many areas of engineering and the sciences and thus the mathematical example here can provide guidance to a wide range of users with models consisting of partial differential equations.

1 Introduction

The study of chemical reaction/diffusion phenomena includes the possibility of multi-scale situations in which one (or more) of the reactions involved, e.g.,

$$\mathrm{A} + \mathrm{B} \xrightarrow{\lambda} \mathrm{C} \tag{1}$$

is much faster than other reaction as well as faster than the time scale of diffusion, i.e., $\lambda \gg 1$. As an initial transient, one then expects the reaction to progress as if in isolation in any small subregion, faster than any diffusive mass transport, until the local concentration of the minority component becomes negligible. Thus, one would develop a regime of (approximately) pure A-regions and B-regions — a free boundary problem. This insight goes back at least to Nernst [4] over a century ago. This modeling ansatz, together with the available computational results and recently developed theory [8] suggests a first simplification: that a limit solution exists as $\lambda \to \infty$ so results for any large λ would be well-approximated by those for any other or by the limiting idealization: "$\lambda = \infty$."

In this paper, continuing the analysis of [9], we consider a model problem in which some reaction $2\mathrm{A} + \mathrm{B} \to (*)$ is given by a reaction pathway involving an intermediate compound C

$$\begin{aligned} \mathrm{A} + \mathrm{B} &\xrightarrow{\lambda} \mathrm{C}, \\ \mathrm{A} + \mathrm{C} &\xrightarrow{\mu} (*) \end{aligned}$$

with $\mu = 1$, but with the first reaction very fast, i.e., $\lambda \gg 1$, as in (1). With u, v, w denoting the concentrations of A,B,C, respectively and simplifying by assuming a common diffusion coefficient,

standard chemical kinetics gives the system

$$\left.\begin{array}{rcrcrcl} u_t &=& \Delta u &-& \lambda uv &-& uw \\ v_t &=& \Delta v &-& \lambda uv & & \\ w_t &=& \Delta w &+& \lambda uv &-& uw \end{array}\right\} \text{ in } \Omega \text{ for } 0 < t \le t_{\text{fin}} \tag{2}$$

with suitable initial data for the three species u, v, w. Note that for existence of an idealized limit solution one must have boundedness, in some sense, of the fast reaction term $q = \lambda uv$ whence, on dividing by $\lambda \to \infty$, one gets $uv \to 0$, i.e., one would have

$$uv \equiv 0. \tag{3}$$

for the idealization. This just corresponds to the separation suggested above into A- and B-regions where u or v, but not both, would be nonzero.

The boundary conditions we consider for this model problem are

$$\begin{array}{c} u = \alpha \text{ on } \Gamma_{\text{A}}, \qquad v = \beta \text{ on } \Gamma_{\text{B}}, \\ \text{no-flux for } u, v, w \text{ otherwise on } \partial\Omega, \end{array} \tag{4}$$

where Γ_{A} and Γ_{B} are disjoint, nontrivial parts of the boundary $\partial\Omega$ and α, β are suitable positive bounded functions, e.g., constants. Note that this setup represents the case of unlimited supply of A at one part of the boundary of the domain Ω and an unlimited supply of B at another; once produced by the fast reaction (1), the intermediate compound C cannot escape from Ω, although it may be consumed by the second, slower, reaction.

Bearing in mind Hamming's dictum that, "The purpose of computation is insight, not numbers," what we seek to learn from computational simulation is

1. support for the conjectured convergence as $\lambda \to \infty$
2. support for the conjectured convergence to steady state as $t \to \infty$
3. a view of the interesting topological changes (A- and B-regions merging, splitting, and disappearing)
4. a view of the fine structure of the interface

Some of this is now already proven, specifically the convergence as $\lambda \to \infty$ was a conjecture at the initiation of this computation, but is now theoretically justified [8]. However, as noted therein, most is in support of conjectures which still remain theoretically open and of analysis yet to be undertaken.

For analysis or simulation of this situation a principal concern lies in treatment of the interfacial boundaries separating these regions. Computationally, the greatest difficulty in handling the system (2) is the occurrence in each of the equations of the fast reaction term $q = \lambda uv$ with $\lambda \gg 1$. Chemical modeling and the rigorous analysis [2] available for the 1-dimensional steady state system show in that setting that this pointwise reaction rate is very large where relevant, in a narrowly concentrated reaction zone: we expect this to be negligible where one or the other of A, B dominate, but to be significant where the diffusion transports these reactants to meet each other at the interface. For computation one needs fairly accurate determination of the integral of q so one must adequately resolve the q profile — a 'spike' in 1-D, a 'wall' in 2-D, etc. — with the further difficulty that the location of this 'spike' is not known a priori. Singular perturbation analysis for the corresponding stationary problem in one spatial dimension is available in [2, 7], and shows the existence of an

internal layer of width $\mathcal{O}(\varepsilon)$ and height $1/\mathcal{O}(\varepsilon)$ with $\varepsilon = \lambda^{-1/3}$; thus, q can be expected to be locally large with sharp gradients near the interface, even though the respective concentrations u, v would each be quite small there. It is conjectured that (away from topology changes) the q-profile transverse to the interface is the same for the n-dimensional time-dependent problem as was shown in [2] for the 1-dimensional steady state problem, i.e., as noted above. In one spatial dimension this has already been considered computationally in [3, 9] and the numerical studies in [9] allow us to conclude that the moving, sharp, internal layer in the transient problem has the same scaling as in the stationary problem.

If one works with the reaction/diffusion system (2) with $\lambda \gg 1$, then one must be able to resolve this fine structure of the q-profile. While this fine structure disappears in the free boundary formulation with a dynamically evolving sharp boundary, such a formulation requires derivation of equations of motion for the interfaces [6] and computational implementation requires a sophisticated information structure for tracking the interface, especially through possible changes of topology. Each of these alternatives would require a fairly fine mesh locally and, especially for two- and three-dimensional settings, this rapidly becomes prohibitive. As is demonstrated in [1] for values $\lambda = 10^3$, 10^6, and particularly for $\lambda = 10^9$, the increasingly steep gradients in the model make the numerical simulations of (2) together with (4) with large values of the fast reaction parameter λ both challenging and costly.

Two-Component Model: What saves the situation is the availability of a reformulation of the system (2) introduced in [6] to avoid working directly with the difficult term $q = \lambda uv$. We begin by noting formally that subtracting the second equation of (2) from the first or adding it to the third will each eliminate this rapid reaction term, giving the pair of equations

$$\begin{array}{rcl} (u-v)_t &=& \Delta(u-v) - uw, \\ (v+w)_t &=& \Delta(v+w) - uw, \end{array} \tag{5}$$

suggesting the use of

$$\begin{array}{rcl} u_1 &:=& u - v, \\ u_2 &:=& v + w. \end{array} \tag{6}$$

as new unknown variables. Since λ no longer appears in (6), we can apply this to the limit problem with $\lambda = \infty$. Since we then have (3) and the components u, v, w are necessarily nonnegative, it is possible to recover all three from the two variables u_1, u_2 by

$$\begin{array}{rcl} u &=& u_1^+ := \max\{u_1, 0\}, \\ v &=& -u_1^- := -\min\{u_1, 0\}, \\ w &=& u_2 + u_1^- = u_2 + \min\{u_1, 0\}, \end{array} \tag{7}$$

effectively inverting (6). This gives $uw = u_1^+ u_2$ so the coupled system (5) becomes

$$\left.\begin{array}{rcl} u_{1,t} &=& \Delta u_1 - u_1^+ u_2 \\ u_{2,t} &=& \Delta u_1 - u_1^+ u_2 \end{array}\right\} \text{ in } \Omega \text{ for } 0 < t \le t_{\text{fin}}. \tag{8}$$

The key to the reformulation, following [6], is that the boundary conditions (4) also permit a completely self-contained version for (8): on Γ_A we have $u = \alpha > 0$ so $v = 0$ there while at Γ_B we have locally $u_1 = -v$ so $u_2 + u_1 = w$ there with $w_\nu = 0$. Hence, from (4) we obtain for (8) the boundary conditions

$$\begin{array}{ll} u_1 = \alpha \text{ on } \Gamma_{\text{A}}, \quad u_1 = -\beta \text{ on } \Gamma_{\text{B}}, & \text{no-flux for } u_1 \text{ otherwise on } \partial\Omega, \\ \text{no-flux for } u_2 \text{ on } \Gamma_A, & \dfrac{\partial u_2}{\partial \nu} = -\dfrac{\partial u_1}{\partial \nu} \text{ otherwise on } \partial\Omega, \end{array} \tag{9}$$

and the two-component system (8)–(9) is self-contained and, given suitable initial data, provides, with (7), a computational alternative to the original three-component system (2) and (4).

This new formulation promises to be significantly more computationally efficient, since, without the appearance of q, it lacks the sharp internal layers present in the three species model: these internal layers remain implicit in the two component model, showing up as discontinuities in spatial derivatives for the original species, but, rather than a term λuv in the equation with $\lambda \approx \infty$, the computation need not deal with any discontinuity worse than the extraction of a positive part in obtaining the (slow) reaction term $u_1^+ u_2$.

We do note, however, an awkwardness of the two component model, involving a boundary condition which couples the fluxes (derivatives) of the solution components. This and other properties of the desired simulations make this an excellent example to bring out the benefits of several features of COMSOL Multiphysics. For instance, COMSOL can handle the boundary coupling of the solution components, and its General Form of the PDEs offers a convenient way to enter the problem into the software and take full advantage of an automatically differentiated Jacobian.

Our goals in this paper are to present evidence that the two component model, first presented in [6], is easily implementable, does indeed gives equivalent results to the three species model in one and in two spatial dimensions, and is significantly computationally superior — showing that in fact the simulations with the two component model are faster than the original three species ones with any finite λ. Thus, this paper shows an interesting case of interaction between analysis and simulations on an application problem with significant mathematical challenges in the form of asymptotic limits and moving interfaces.

In this paper, we consider this computation for a 1-dimensional setting

$$\Omega = (0,1) \text{ with } \Gamma_A = \{0\} \text{ and } \Gamma_B = \{1\} \tag{10}$$

and for a 2-dimensional setting

$$\Omega = (0,1) \times (0,1) \text{ with } \Gamma_A = \{0\} \times (0,1) \text{ and } \Gamma_B = \{1\} \times (0,1). \tag{11}$$

For the boundary data we take α, β in $u = \alpha$ on Γ_A and $v = \beta$ on Γ_B to be the same in (10) and (11) to enable comparisons. Also, in this configuration, the steady state of the two-dimensional problem will be the same as that of the one-dimensional one for each y value. For the initial data we take u, v to already have disjoint supports and to be consistent with the boundary conditions.

More details on the numerical method are specified in Section 2. Sections 3 and 4 present the results for numerical studies in one and in two spatial dimensions, respectively. Subsections 3.1 and 4.1 contain graphical comparisons between simulations for the three species model with finite, large λ and the two species model with $\lambda = \infty$; we use the phrase "$\lambda = \infty$" here and in the following as a short hand notation to indicate the use of the two component model. Subsections 3.2 and 4.2 contain accuracy and efficiency comparisons for selected critical solution quantities that demonstrate both the reliability of simulations with the two component model and its superior computational efficiency.

2 Numerical Method

The computations use the state-of-the-art commercial finite element package COMSOL Multiphysics (`www.comsol.com`). Several features of the problem under consideration make it particularly profitable to use this software:

One of the crucial features of the reaction-diffusion equations under consideration are their non-linear reaction terms. Since the physics of this problem relies on a subtle balance between the diffusion effects and these reactions, we wish to represent them as accurately as possible. Thus, we use the General Form of the PDEs in COMSOL in order to enable COMSOL to compute symbolic derivatives of all terms in the PDEs by automatic differentiation for the highest accuracy in the evaluation of the Jacobian in the non-linear solver inside the implicit ODE method. By contrast, the Coefficient Form in COMSOL approximates the derivatives of all terms numerically, which is only suitable for mildly non-linear PDEs.

An interesting feature of the two component problem is the boundary condition $\frac{\partial u_2}{\partial \nu} = -\frac{\partial u_1}{\partial \nu}$ on a portion of $\partial\Omega$ that couples the (derivatives of the) solution components. Many PDE solvers have difficulty handling a boundary condition of this type or do not permit specifying it at all, even for the problem in only one spatial dimension (e.g., Matlab's `pdepe` function). COMSOL has no trouble with this, as all terms in the boundary condition are allowed to include also the dependent variables and their derivatives. We also note that the 2-D problem is not any harder to implement than the 1-D problem in a software like COMSOL.

We use COMSOL coupled with Matlab. With this, a script is written so that each computation would be using the same solver parameters inside COMSOL. This is crucial to ensure reproducibility of all results and useful to facilitate parameter studies, by that with respect to problem parameters or with respect to numerical parameters. Scripting is also very useful in cases such as the two component model, where significant post-processing such as back-transforming to the original variables by (7) is necessary. But we use the flexibility of scripting even more for the problems under consideration here. COMSOL easily allows the user to specify functional expressions for boundary and initial conditions. This is suitable for the boundary conditions here, and changing one of them is possible in the script. But the initial conditions desired here will have a fairly complicated structure inside the domain, and we moreover wish to make changing the initial condition convenient. Using functional expressions in the script becomes too cumbersome and potentially limiting. Therefore, use merely specify the names of functions for the initial conditions in the script, which enables the user to write these functions outside of COMSOL using the full capabilities of high-level programming languages such as if-statements, for-loops, and vector operations on the coordinate values.

In 1-D, we use linear Lagrange finite elements on a uniform mesh with N intervale elements. In 2-D, we use linear Lagrange finite elements on a uniform quadrilateral mesh with $N \times N$ elements. Note that the defaults in COMSOL for 2-D are an unstructured mesh and quadratic Lagrange finite elements. By using a structured, uniform, quadrilateral mesh, we avoid any incidental biasing that might affect the solution as result of, e.g., from element boundaries being at different random non-horizontal or non-vertical angles or similar properties inherent to an unstructured mesh. Because we will be interested in the exact location of interfaces later, it is important to avoid any such biasing for this problem. Because the solutions can have jump discontinuities at the initial time and will have discontinuities in their derivatives at the reaction interfaces [6], we use the lowest order Lagrange finite elements available.

COMSOL offers several ODE solvers, all of which use implicit time stepping, which is a necessity for efficiently solving PDEs of parabolic type such as reaction-diffusion equations. Since we desire to compute to a large final time approaching steady state, sophisticated ODE solver features such

as automatic time stepping and method order selection up to a high order are vital. This becomes particularly clear if you consider that we have to expect steep initial gradients from the chosen initial conditions, which requires small time steps and low ODE method orders, but that we wish to use large time steps and high method orders, when the solution is smooth in its approach to steady state. In 1-D, we use the default ODE solver BDF-IDA with its default tolerances for the local error control, namely, a relative tolerance of 10^{-2} and an absolute tolerance of 10^{-3}. In 2-D, the ODE solver used is BDF-DASPK. We use a relative tolerance of 10^{-3} and an absolute tolerance of 10^{-6} for the local error control in the ODE solver, which is slightly tighter than COMSOL defaults. We experimented with coarser as well as tighter tolerances. At coarser tolerances, some features of the solution are not as clear. Tighter tolerances confirm the results obtained for the tolerances used, thus these are the most effective tolerances to use. The increased numerical difficulty associated with large λ values materializes also in other ways, e.g., that the ODE solver broke down for the coarse 64×64 mesh at an intermediate time and we were forced to use a coarser ODE tolerance for this case. We choose the ODE solver BDF-DASPK because BDF-IDA, the default solver, has trouble converging at the initial conditions for the case of large λ values. The ODE solver Generalized Alpha was also experimented with, but this too had some difficulties.

In 1-D, we use the default linear solver UMFPACK. In 2-D, we use the linear solver PARDISO, since it is supposed to profit most from the multi-threading available on the multi-core processors used [5].

3 Numerical Studies in One Spatial Dimension

This section considers the chemical reaction-diffusion problem in one spatial dimension given by (2), (4), and (10) in the domain $\Omega = (0, 1)$ and $0 \le t \le 20$.

In the first subsection, we present qualitative evidence that over time both the three species model and two component model give equivalent results. In the second subsection, we give a quantitative numerical comparison of the two models based on the grid size used in the computations.

3.1 Qualitative Graphical Comparisons

We let the initial conditions be of the continuous form

$$u_{\text{ini}}(x) = \begin{cases} 4(0.25 - x)\,\alpha, & \text{if } 0.00 \le x \le 0.25, \\ 0, & \text{if } 0.25 \le x \le 0.50, \\ 64(0.50 - x)(x - 0.75)\,\gamma, & \text{if } 0.50 \le x \le 0.75, \\ 0, & \text{if } 0.75 \le x \le 1.00, \end{cases}$$

$$v_{\text{ini}}(x) = \begin{cases} 0, & \text{if } 0.00 \le x \le 0.25, \\ 64(0.25 - x)(x - 0.50)\,\delta, & \text{if } 0.25 \le x \le 0.50, \\ 0, & \text{if } 0.50 \le x \le 0.75, \\ 4(x - 0.75)\,\beta, & \text{if } 0.75 \le x \le 1.00, \end{cases}$$

$$w_{\text{ini}}(x) = 0.$$

Here we pick $\alpha = 1.6$, $\beta = 0.8$, and $\gamma = \delta = 0.25$. Notice that the initial conditions satisfy the boundary conditions. Also, $u_{\text{ini}} v_{\text{ini}} \equiv 0$ in Ω with three initial reaction interfaces located at $x = 0.25$, 0.50, and 0.75; compare this with the fact that the steady state solution has only one interface located at $x^*_{\text{SS}} \approx 0.6$. For more studies of these types of initial conditions, see [9].

The two component model in one spatial dimension given by (8), (9), and (10). Notice that for the case of (10) with $\Omega = (0, 1)$ and $\Gamma_{\text{A}} = \{0\}$, the coupling condition from (9) becomes concretely $u_{2,x} = -u_{1,x}$ at $x = 1$. Using the transformations for the two component model with $\lambda = \infty$, its initial conditions are

$$u_{1,\text{ini}}(x) = \begin{cases} 4(0.25 - x)\,\alpha, & \text{if } 0.00 \le x \le 0.25, \\ -64(0.25 - x)(x - 0.50)\,\delta, & \text{if } 0.25 \le x \le 0.50, \\ 64(0.50 - x)(x - 0.75)\,\gamma, & \text{if } 0.50 \le x \le 0.75, \\ 4(0.75 - x)\,\beta, & \text{if } 0.75 \le x \le 1.00, \end{cases}$$

$$u_{2,\text{ini}}(x) = \begin{cases} 0, & \text{if } 0.00 \le x \le 0.25, \\ 64(0.25 - x)(x - 0.50)\,\delta, & \text{if } 0.25 \le x \le 0.50, \\ 0, & \text{if } 0.50 \le x \le 0.75, \\ 4(x - 0.75)\,\beta, & \text{if } 0.75 \le x \le 1.00, \end{cases}$$

Notice that the initial conditions satisfy the boundary conditions, including the coupling condition $u_{2,x} = -u_{1,x}$ at $x = 1$.

The following figures are simulations for the solutions and interfaces at the final times for the three species model with $\lambda = 10^6$ in Figure 1 and $\lambda = 10^9$ in Figure 2 and the two component model in Figure 3. The waterfall plots in (a), (b), and (c) in each figure represent the concentrations of $u(x, t)$, $v(x, t)$, and $w(x, t)$, respectively, over (x, t). The horizontal axis in both (d) and (e) is the central portion of the domain $\Omega = (0, 1)$ including the range 0.25 to 0.75, where the interfaces are located in the initial condition. Subplot (d) shows the movement of the interface over the entire period time $0 \le t \le 20$, and subplot (e) is a zoom with respect to time that focuses on the initial

period for $0 \leq t \leq 0.1$. In both (d) and (e), we can see the interfaces begin at the locations $x = 0.25$, $x = 0.5$, and $x = 0.75$ at time $t = 0$ (i.e., bottom of each plot), as specified by the initial condition. They coalesce to a single interface very quickly, the specific dynamics of which are more clearly visible in the zoomed image in subplot (e), which confirms the coalescing of two interfaces and the movement of the third to the steady state position around $x^* \approx 0.6$.

Comparing Figure 1 and Figure 2, we see very similar movement in the interfaces. The waterfall plots also appear incredibly similar. At $t = 0$, see the non-zero initial values for u and v in one interior region each, which dissipates rapidly during an initial transient phase of the evolution. This offers further evidence that as we take $\lambda \to \infty$, we will obtain similar results. For more studies of initial conditions of this form, see [3]. Simulation results for solutions and interface for the two component model with $\lambda = \infty$ are shown as Figure 3. Comparing them with Figures 1 and 2 obtained by the three species model, we observe that Figure 3 is very similar. Thus we can confirm qualitatively that the two component model is reasonable and that there is an asymptotic consistency in the concentrations and interfaces when $\lambda \to \infty$.

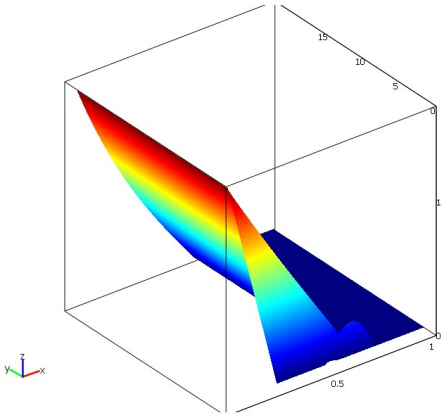

(a) u vs. (x, t)

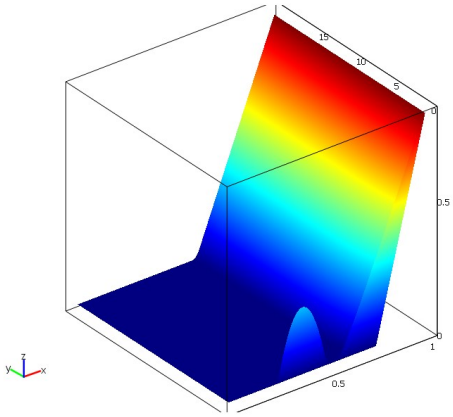

(b) v vs. (x, t)

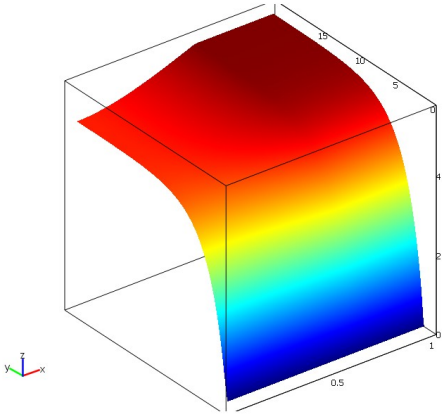

(c) w vs. (x, t)

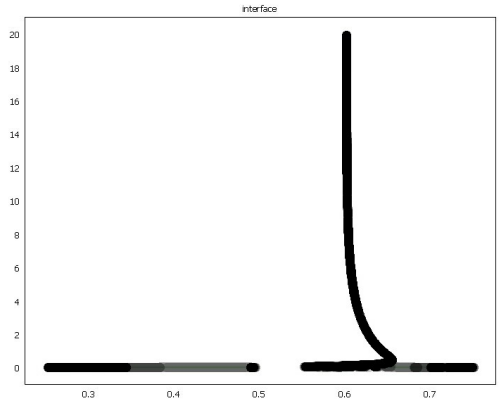

(d) interface vs. (x, t)

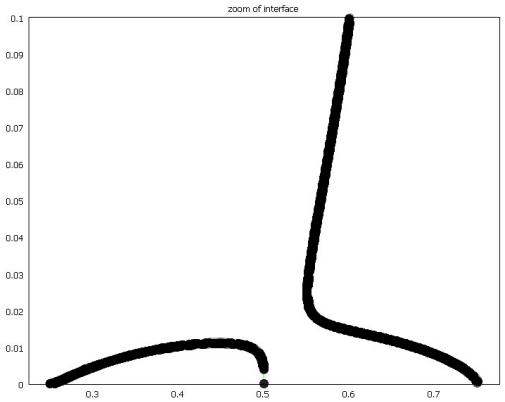

(e) zoomed interface vs. (x, t)

Figure 1: Simulations of the three species model with $\lambda = 10^6$.

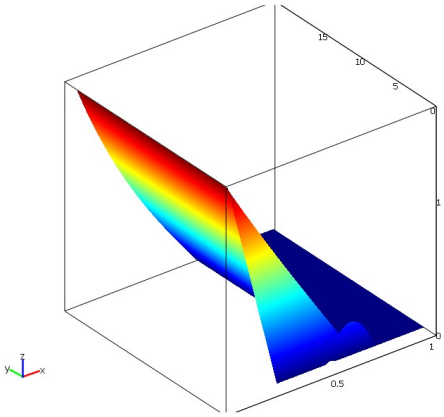

(a) u vs. (x, t)

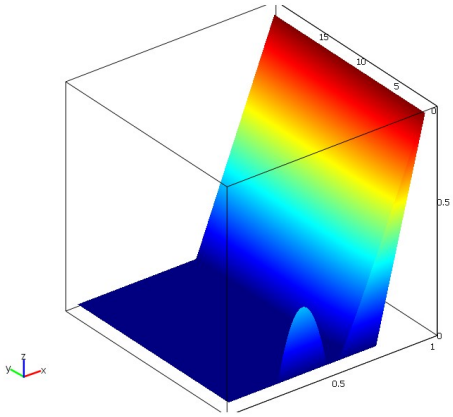

(b) v vs. (x, t)

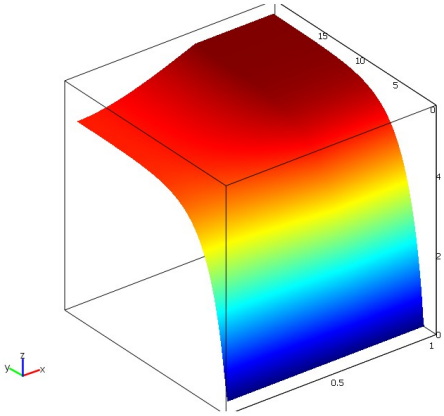

(c) w vs. (x, t)

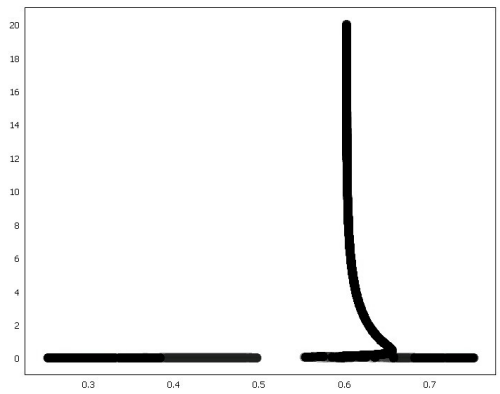

(d) interface vs. (x, t)

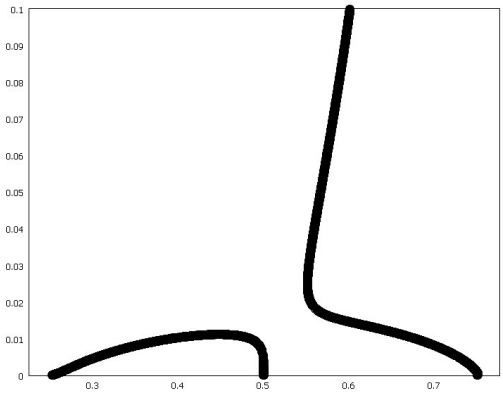

(e) zoomed interface vs. (x, t)

Figure 2: Simulations of the three species model with $\lambda = 10^9$.

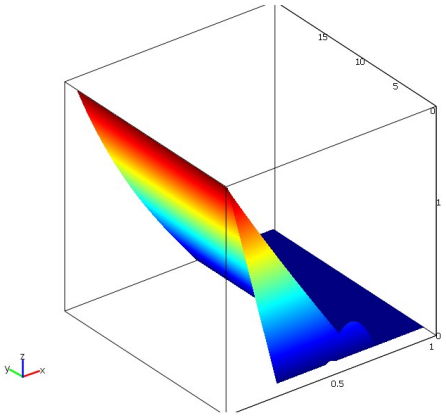

(a) u vs. (x, t)

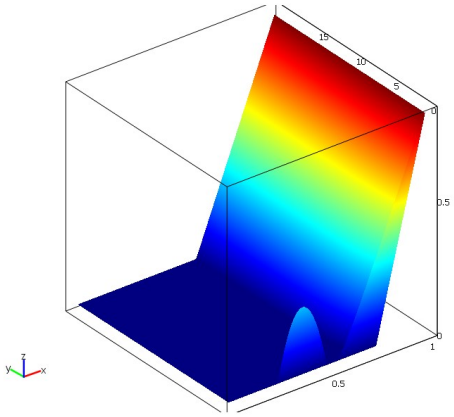

(b) v vs. (x, t)

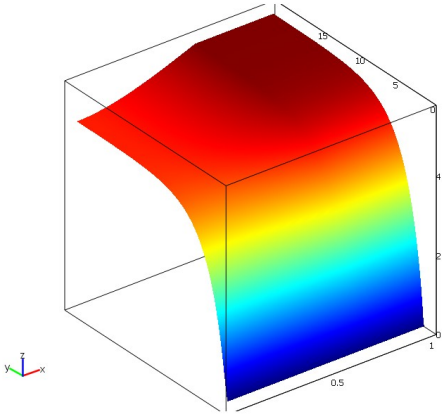

(c) w vs. (x, t)

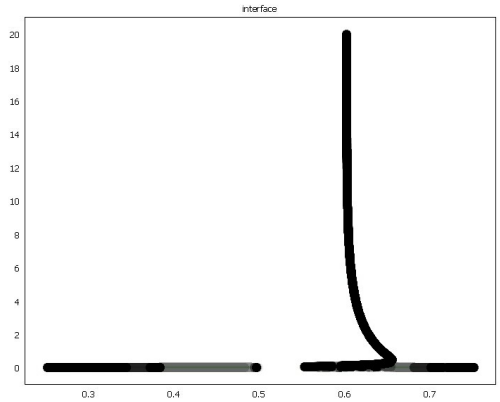

(d) interface vs. (x, t)

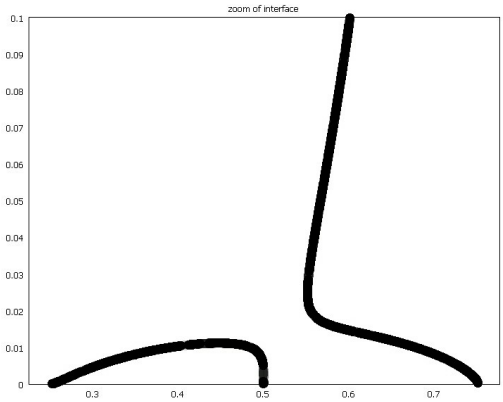

(e) zoomed interface vs. (x, t)

Figure 3: Simulations of the two component model with $\lambda = \infty$.

3.2 Quantitative Accuracy and Efficiency Comparisons

As our goal is to verify that the two component model is indeed numerically superior to the three species model, we need to validate this claim quantitatively. So in order to do this, we compare two measures of accuracy and two measures of efficiency, over a variety of finite element meshes with N elements. The first measure of accuracy is the closeness of the location of the interface x^* at the final time of the transient problem to that of the actual steady state interface, which is $x^*_{\rm SS} = 0.601806640625$; this value was obtained by simulations of the three species steady state problem on a high resolution mesh with $N = 8192$ and will be considered the true value for x^* in the following. The second measure of accuracy is the time $t_{\rm co}$ required for all three interfaces from the initial conditions above to coalesce into a single interface. Our measures of efficiency are fairly standard; we consider the number of time steps taken by the ODE solver and the total computation time in seconds (s) taken by COMSOL. The data are summarized in the Table 1.

Table 1: Summary of accuracy and efficiency data for simulations (a) of the three species model with $\lambda = 10^6$, (b) of the three species model with $\lambda = 10^9$, and (c) of the two component model with $\lambda = \infty$.

(a) $\lambda = 10^6$				
N	x^*	$t_{\rm co}$	steps	time (s)
128	0.6015625000000	0.0112182294336213	552	8.860
256	0.6015625000000	0.0112334958601752	465	8.312
512	0.6015625000000	0.0111585034464878	524	10.937
1024	0.6015625000000	0.0111818581067737	513	14.219
2048	0.6020507812500	0.0111876013688106	516	18.359
4096	0.6018066406250	0.0112575048005446	524	26.577
(b) $\lambda = 10^9$				
N	x^*	$t_{\rm co}$	steps	time (s)
128	0.6015625000000	0.0111557076128189	1,079	21.468
256	0.6015625000000	0.0111679246895327	1,340	30.062
512	0.6015625000000	0.0111606912657899	1,483	37.843
1024	0.6015625000000	0.0111663934378730	1,461	55.170
2048		out of memory		
4096		out of memory		
(c) $\lambda = \infty$				
N	x^*	$t_{\rm co}$	steps	time (s)
128	0.6015625000000	0.0112999361134998	318	2.891
256	0.6015625000000	0.0112894187459985	297	2.891
512	0.6015625000000	0.0112529236404671	285	3.031
1024	0.6015625000000	0.0112219691583494	275	3.516
2048	0.6020507812500	0.0112475828870521	277	4.985
4096	0.6018066406250	0.0112277130521462	276	8.515

As for accuracy, the table shows that for corresponding numbers of elements N, the location of the interface at the final time is the same for each of $\lambda = 10^6,\ 10^9$, and ∞. The location x^* changes slightly with increasing mesh resolution, because finer meshes have mesh points that are closer to the true value of x^* than coarser meshes. But in all cases of N, the location determined for x^* is the best or one of the equivalently best values possible on that mesh. Therefore, we conclude that

the values for x^* are as accurate as possible for each mesh size. Further, for each of $\lambda = 10^6$, 10^9, and ∞, the coalescing times t_{co} are very close for different numbers of mesh resolutions. Further, for a fixed grid size, the value of t_{co} is similar across all of $\lambda = 10^6,\ 10^9$, and ∞. Thus we confirm that not only does it look like the two component model is giving similar results to the three species model for large λ, but indeed it is giving essentially identical results.

Now from a numerical standpoint, the interesting data are in the last two columns of Table 1. As we can see, the number of time steps required for a single λ over different grid sizes does not vary considerably. However, as we increase λ, the number of time steps required to complete the computation increases considerably. We see that for $\lambda = 10^9$, the solver could not even finish the computations for $N = 2048$ and 4096 on the computer used, because the amount of data per time step is so large combined with the large number of time steps required. Now, when we consider the $\lambda = \infty$ case, we see immediately that both the number of steps taken by the solver and the amount of time are much smaller. Thus, now that we have confirmed that the two component model is both accurate and efficient, this gives strong motivation for using the two component model whenever possible.

4 Numerical Studies in Two Spatial Dimensions

This section considers the chemical reaction-diffusion problem in two spatial dimension given by (2), (4), and (11) in the domain $\Omega = (0,1) \times (0,1)$ and $0 \le t \le 20$.

In the first subsection, we present qualitative evidence that over time both the three species model and two component model give equivalent results. In the second subsection, we give a quantitative numerical comparison of the two models based on the grid size used in the computations.

4.1 Qualitative Graphical Comparisons

We are interested in initial condition functions u_{ini}, v_{ini}, w_{ini} such as shown in Figure 4, where each connected piece of each initial concentration has some constant value. For u in Figure 4 (a), which has an unlimited supply with concentration $\alpha = 1.6$ on $\Gamma_{\text{A}} = \{0\} \times (0,1)$, that is, at $x = 0$, the left hand portion of the domain, all portions of the domain connected to it, as well as the small disjoint disk in the upper right of the domain have values of $u_{\text{ini}} = \alpha$, and $u_{\text{ini}} = 0$ everywhere else. For v in Figure 4 (b), which has an unlimited supply with concentration $\beta = 0.8$ on $\Gamma_{\text{B}} = \{1\} \times (0,1)$, that is, at $x = 1$, the values are $v_{\text{ini}} = 0$ wherever $u_{\text{ini}} > 0$ and $v_{\text{ini}} = \beta$ wherever $u_{\text{ini}} = 0$. We note that by this construction the product $uv \equiv 0$ at $t = 0$. The concentration of w_{ini} is 0 throughout the domain at $t = 0$, as in Figure 4 (c). The fast reaction in the reaction model is restricted to areas of Ω where u and v co-exist. This is only the case along the interface through the domain where u and v come in contact due to diffusion. The locations of these interface lines at $t = 0$ are shown in Figure 4 (d). The interface location is determined numerically as the 0 level of a contour plot of the quantity $u - v$ (with only this one contour level shown in the plot). To allow us a concrete reference to various parts of the domain, we use the following terminology: The portion on the left side of the domain is called the 'body,' and the piece protruding from the body is called the 'head,' which is connected to the body by the 'neck.' The disjointed piece in the upper right of the domain is called the 'disk.' The inspiration for these terms is most evident in the interface plot in Figure 4 (d).

The following eight figures show a comparison of simulation results for the three species model with $\lambda = 10^6$ in Figures 5–8 to the two component model with $\lambda = \infty$ in Figures 9–12; see [1] for additional figures for the three species model with $\lambda = 10^3$ and 10^6. For each simulation, we present the time evolution of u, v, w, and the reaction interface. For the $\lambda = \infty$ case, we back transform u_1 and u_2 into u, v, and w using the transformations $u = \max\{u_1, 0\}$, $v = -\min\{u_1, 0\}$, and $w = u_2 + \min\{u_1, 0\}$ from (7). Each page shows the results at the six times $t = 10^{-4}$, 10^{-3}, 10^{-2}, 10^{-1}, 1, and 20; the initial conditions at $t = 0$ are the same for all simulations and shown in Figure 4. All images were created on a quadrilateral 128×128 mesh using linear Lagrange finite elements.

Specifically, we discuss now the results for $\lambda = 10^6$ in Figures 5, 6, 7, and 8 in detail now. Figures 5, 6, and 7 show the evolution over time of the chemical species u, v, and w, respectively for $\lambda = 10^6$. For example, we see in Figure 5 that over time, u starts from the initial condition in Figure 4 with $u = \alpha = 1.6$ along $x = 0$ and $u = 0$ along $x = 1$. Over a short period of time in the first frame of Figure 5, the sharp edges, which represent a sharp drop off in the amount of chemical, has rounded off as some of the chemical towards the edge reacts with some of the other chemical. As time progresses, we see that the mound has rounded off even more. In the next frame, the mounds have almost completely disappeared. In the last two frames, the chemical has just about settled into its steady state. There is a similar progression in the v images in Figure 6, with $v = 0$ along $x = 0$ and $v = \beta = 0.8$ along $x = 1$. In the w images in Figure 7, we see that the simulation starts with $w = 0$ throughout Ω in Figure 4 and that some $w > 0$ develops first along the reaction

interfaces in the first frame of Figure 7. As time progresses in the next few frames, we see that the areas with less w start to fill in by diffusion from the reaction interface, which is the only place where w is generated. In the following two frames, the peak amount of w still follows the reaction interface. But eventually, the amount of w is highest in the right portion of the domain; this results from the fact that the intermediate species is consumed in a reaction with the first species u, but not the second v, hence it can diffuse from the reaction interface to the right without consumption, while the concentration of w progressively decreases as it diffuses from the reaction interface to the left of the interface. We observe that the growth of w slows down over time towards an apparently finite steady state value; that is a significant observation, for which a rigorous proof is only being developed at present [8]. The most interesting set of images is the collection of interface images in Figure 8. We see here the outlines of the actual reaction interface where $u > 0$ and $v > 0$ meet by diffusion and react rapidly with rate λuv. Recall that at the initial time, these two species do not coexist and mathematically $uv \equiv 0$ at $t = 0$. It is only by diffusion that positive values of both get in contact with each other and this gives the first positive values of the resulting reaction intermediate w. Numerically, we determine the location of the interface as a contour plot with one contour level of value 0 of the difference $u - v$. In the first frame of Figure 8 at $t = 10^{-4}$, we can still clearly see the 'head' connected by the 'neck' to the 'body' on the left portion of the domain and an disjoint 'disk' in the upper right of the domain, like in the initial condition in Figure 4; these outlines have the same shape as the locations of $w > 0$ at the corresponding time in Figure 7. In the next frame, we see the head separating from the body in the left portion of the domain, before it re-attaches to the body in the following frame. Also by the time of the third frame, the disjoint disk has vanished with its supply of the u species completely consumed. The chosen initial conditions give rise to an interesting behavior in that the head separates and then re-attaches to the body in the left portion of the domain, while at the same time the disk disappears. In the last three frames, we see the interface tends to its steady state location at $x^*_{\mathrm{SS}} \approx 0.6$.

Comparing the set of figures for $\lambda = 10^6$ to the corresponding ones for the two component model with $\lambda = \infty$ in Figures 9, 10, 11, and 12, we can clearly see that the progression over time is essentially identical up to the graphical resolution of the studies.

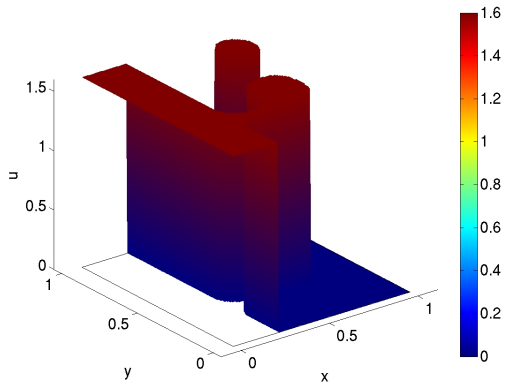

(a) u at $t = 0$

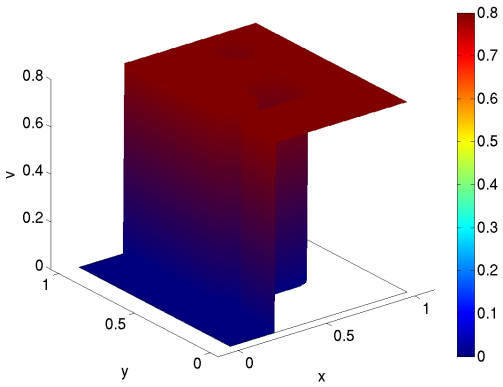

(b) v at $t = 0$

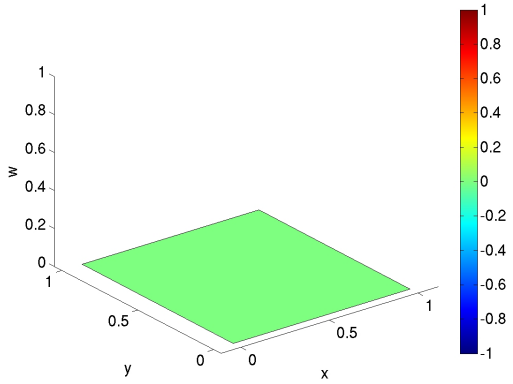

(c) w at $t = 0$

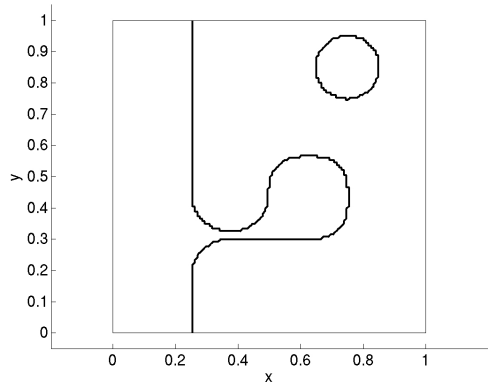

(d) interface at $t = 0$

Figure 4: Initial Condition for u, v, and w, and the initial interface.

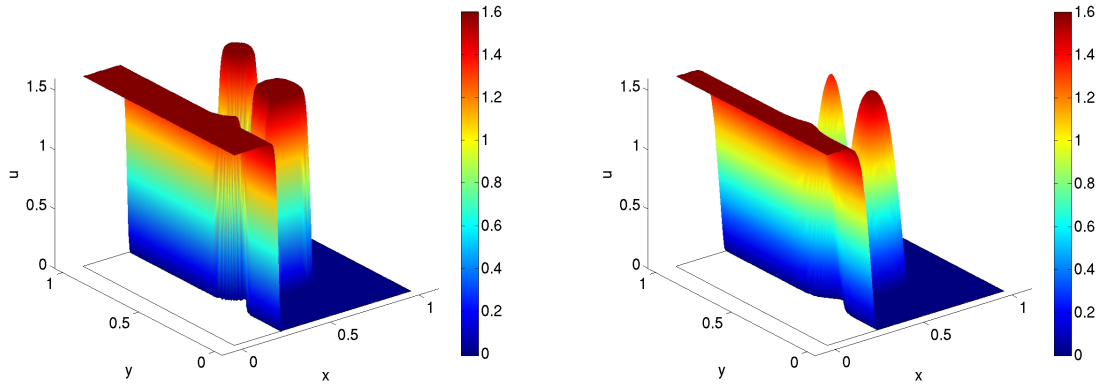

$t = 10^{-4}$

$t = 10^{-3}$

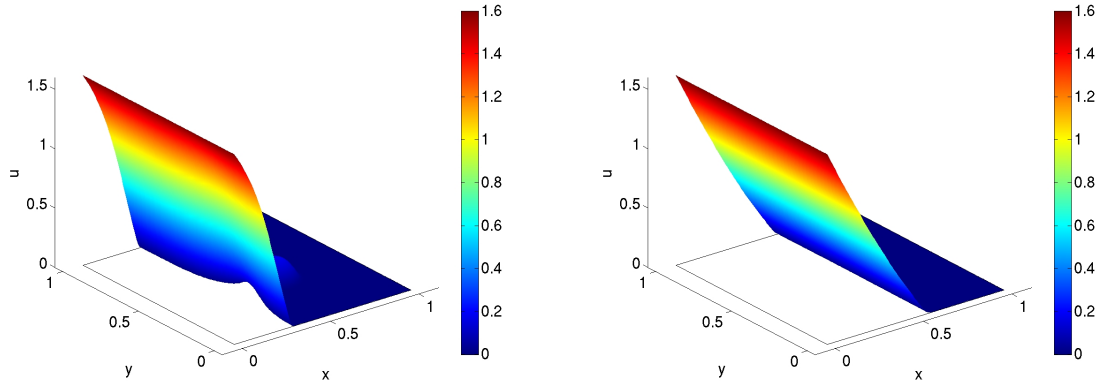

$t = 10^{-2}$

$t = 10^{-1}$

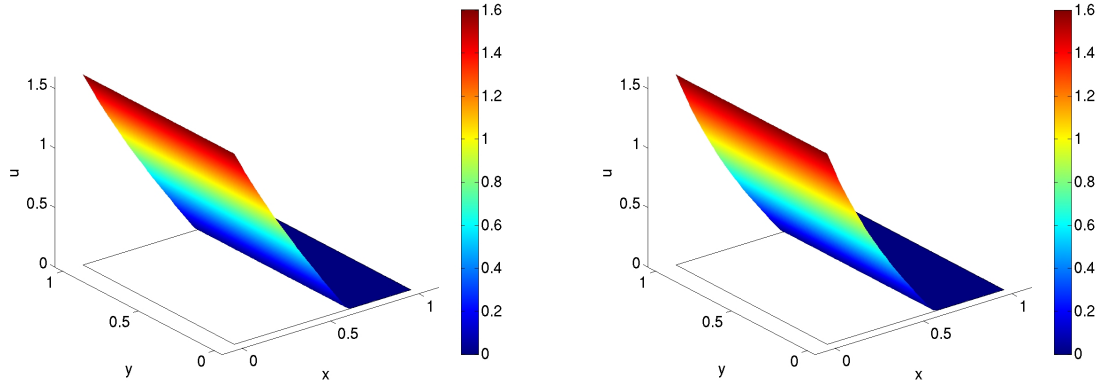

$t = 1$

$t = 20$

Figure 5: $u(x, y, t)$ vs. (x, y) at specified times for $\lambda = 10^6$.

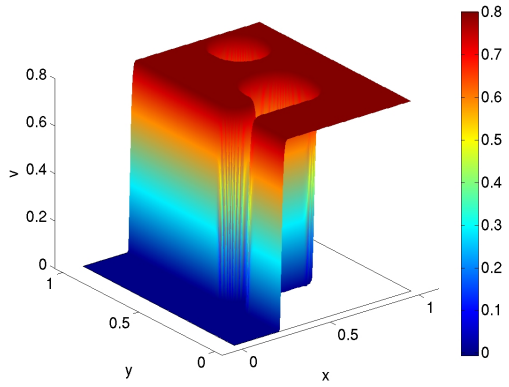

$t = 10^{-4}$

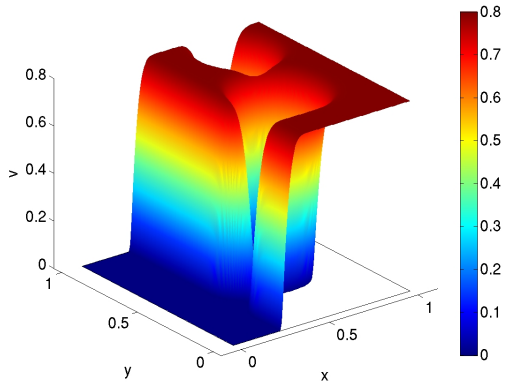

$t = 10^{-3}$

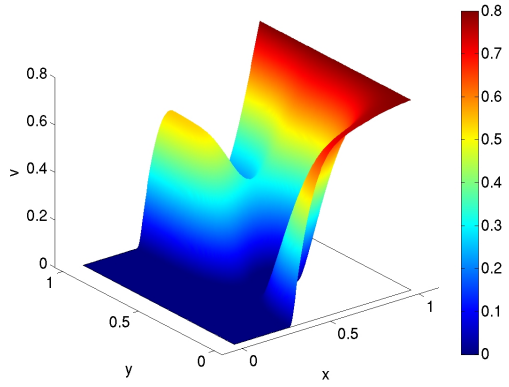

$t = 10^{-2}$

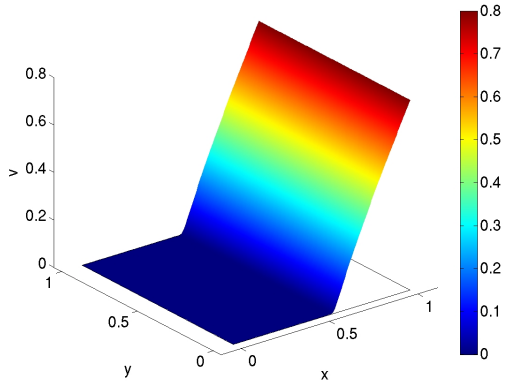

$t = 10^{-1}$

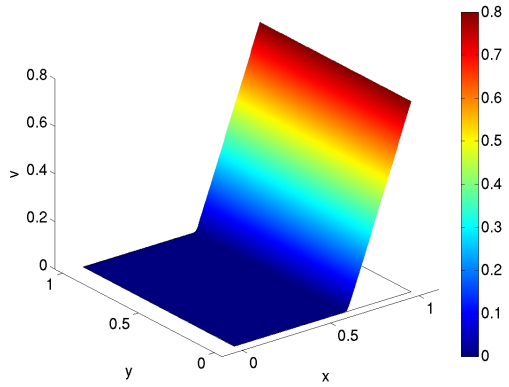

$t = 1$

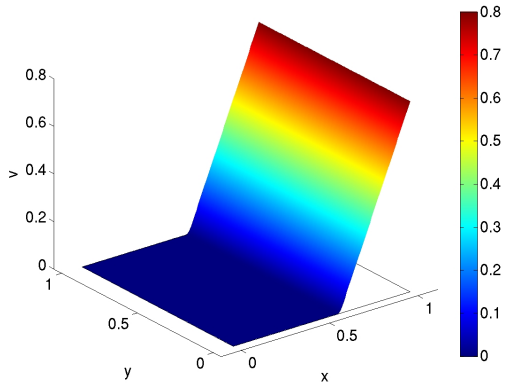

$t = 20$

Figure 6: $v(x, y, t)$ vs. (x, y) at specified times for $\lambda = 10^6$.

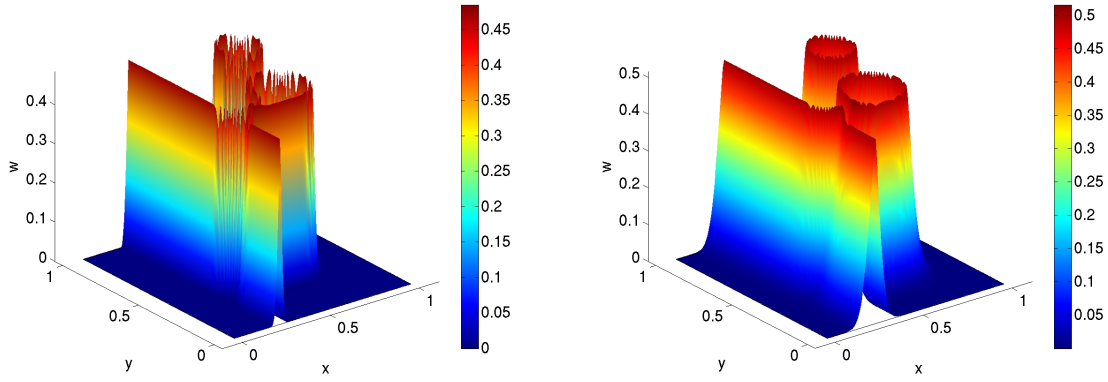

$t = 10^{-4}$ $t = 10^{-3}$

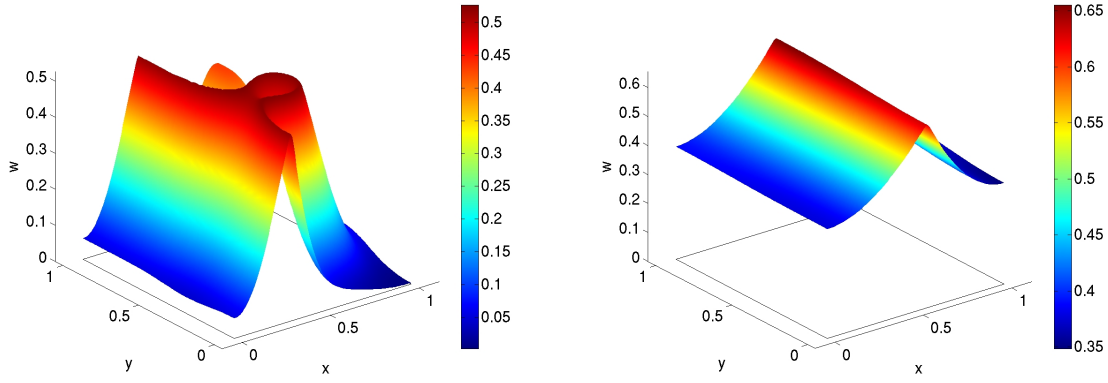

$t = 10^{-2}$ $t = 10^{-1}$

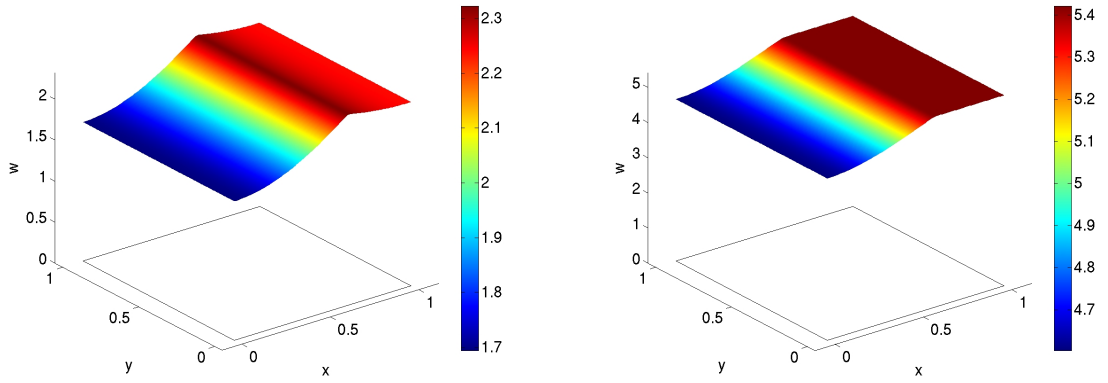

$t = 1$ $t = 20$

Figure 7: $w(x, y, t)$ vs. (x, y) at specified times for $\lambda = 10^6$.

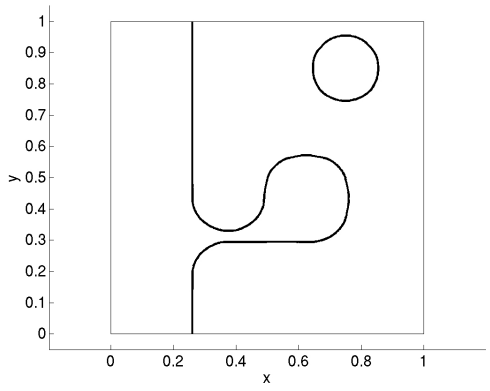

$t = 10^{-4}$

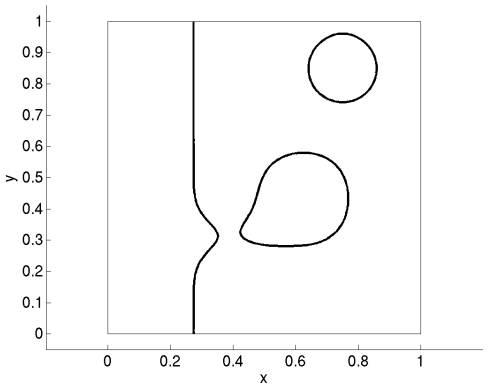

$t = 10^{-3}$

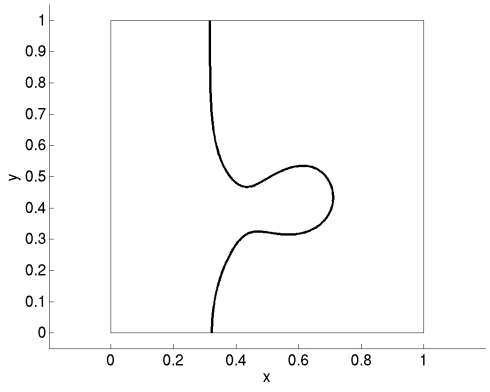

$t = 10^{-2}$

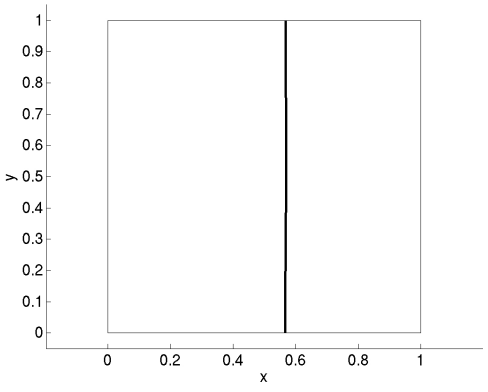

$t = 10^{-1}$

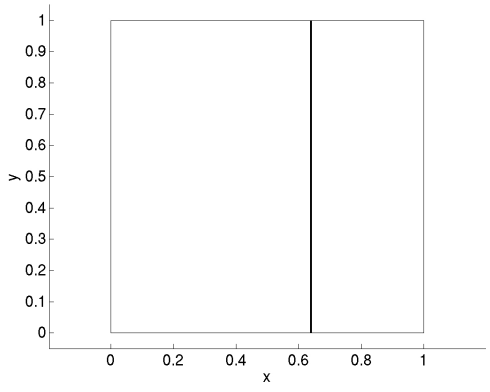

$t = 1$

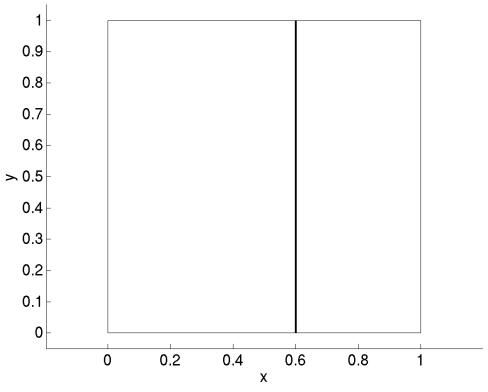

$t = 20$

Figure 8: Interface at specified times for $\lambda = 10^6$.

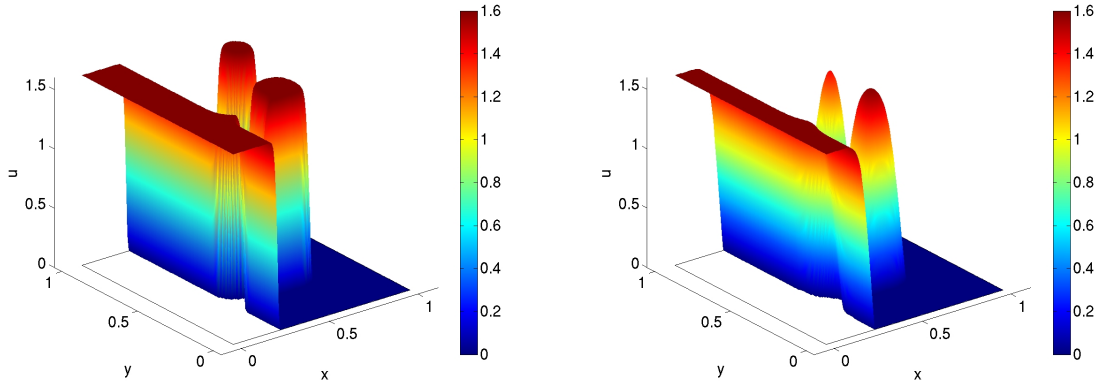

$t = 10^{-4}$ $t = 10^{-3}$

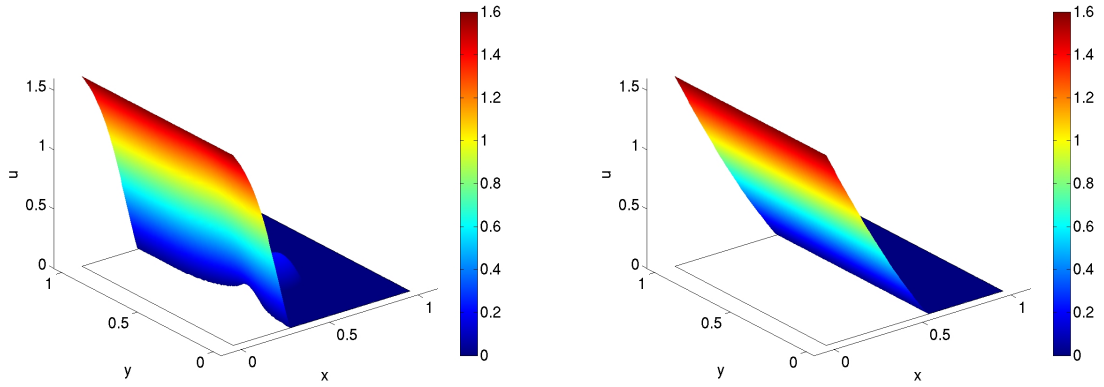

$t = 10^{-2}$ $t = 10^{-1}$

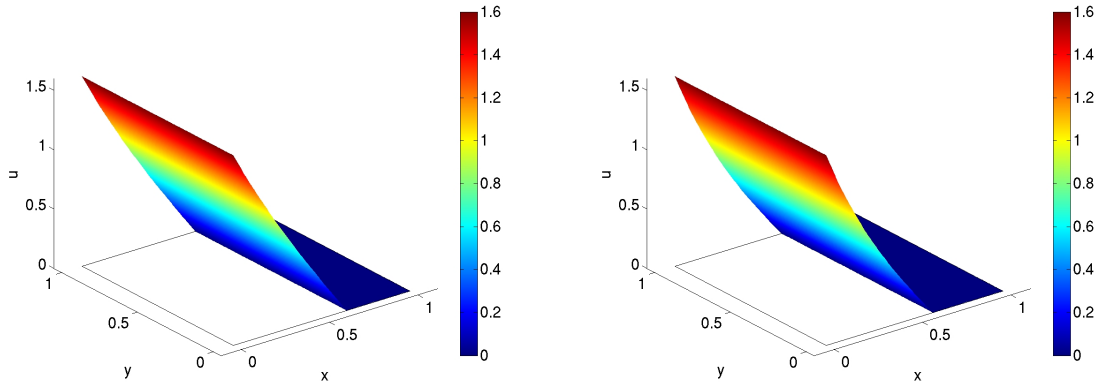

$t = 1$ $t = 20$

Figure 9: $u(x, y, t)$ vs. (x, y) at specified times for $\lambda = \infty$.

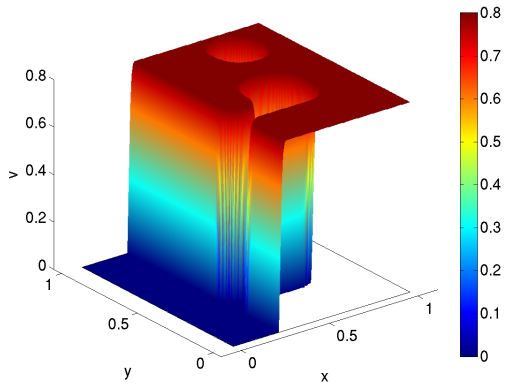

$t = 10^{-4}$

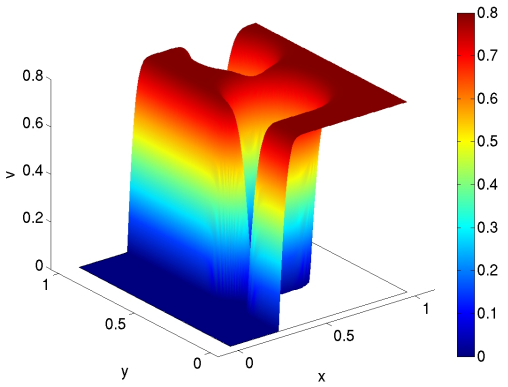

$t = 10^{-3}$

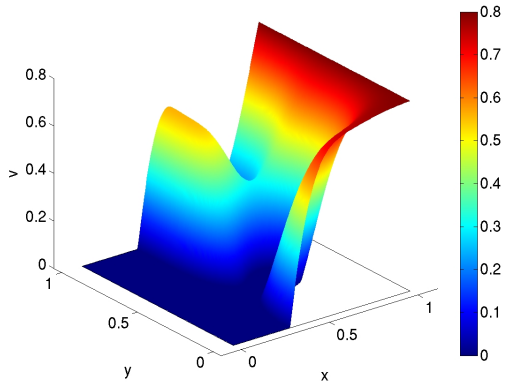

$t = 10^{-2}$

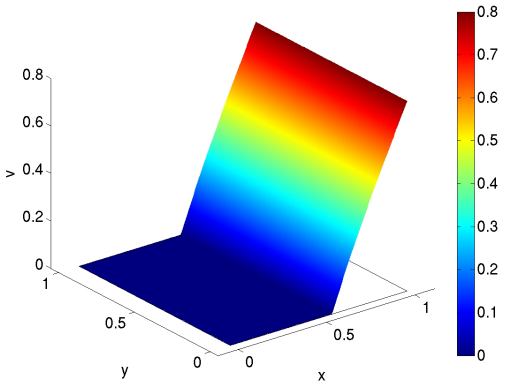

$t = 10^{-1}$

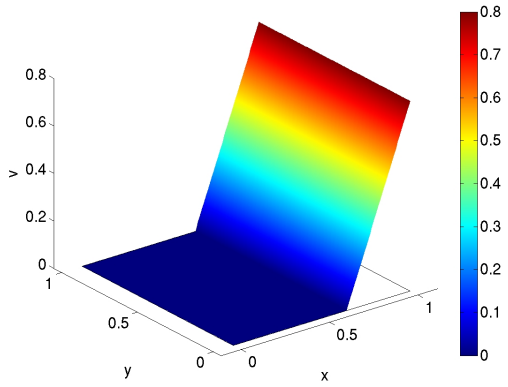

$t = 1$

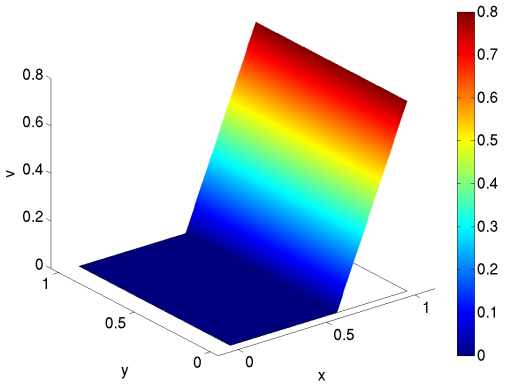

$t = 20$

Figure 10: $v(x, y, t)$ vs. (x, y) at specified times for $\lambda = \infty$.

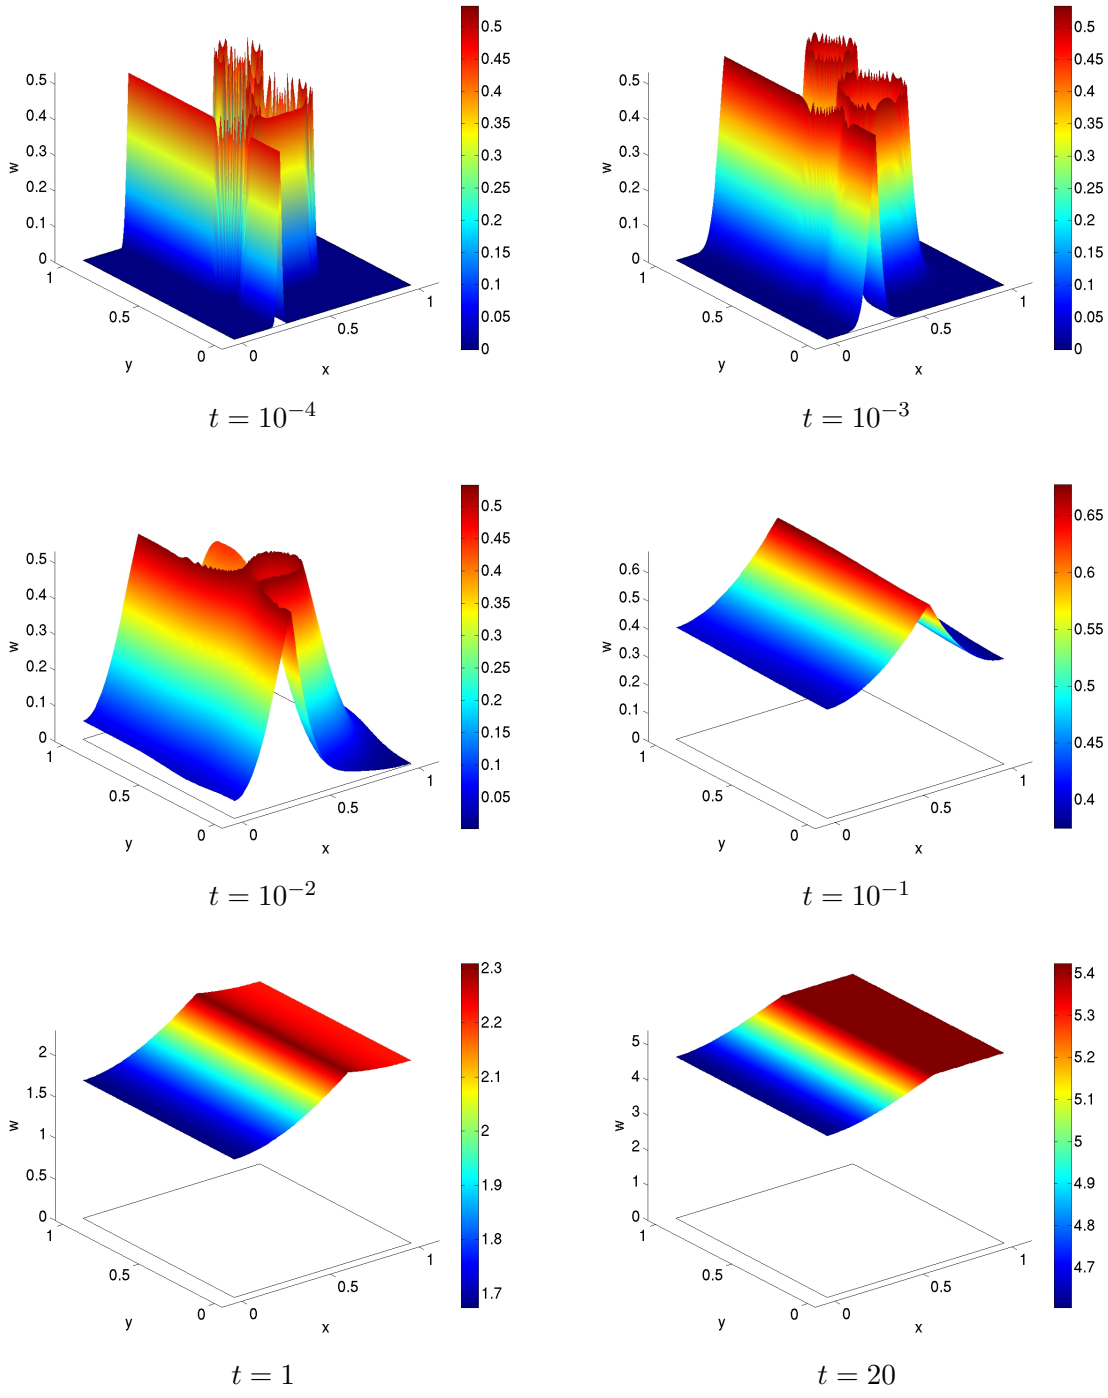

Figure 11: $w(x, y, t)$ vs. (x, y) at specified times for $\lambda = \infty$.

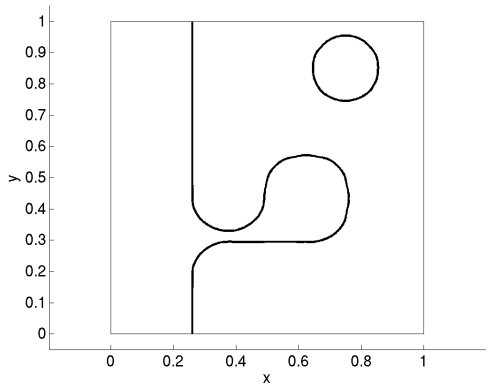

$t = 10^{-4}$

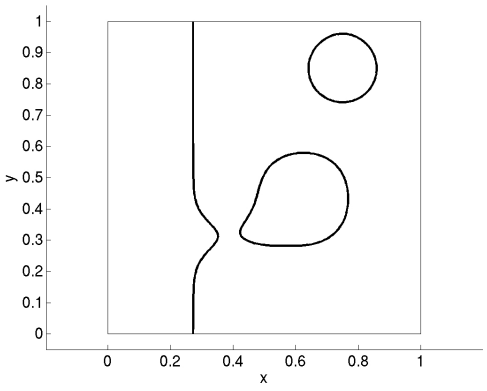

$t = 10^{-3}$

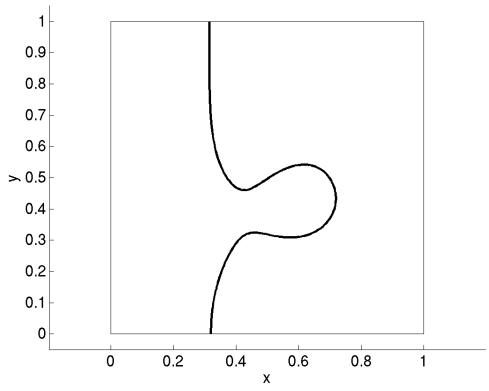

$t = 10^{-2}$

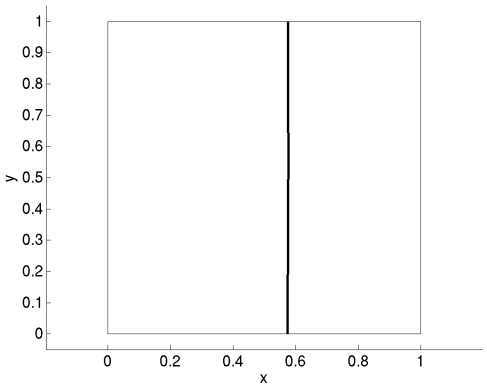

$t = 10^{-1}$

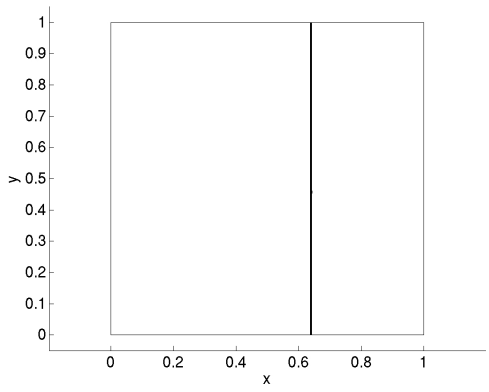

$t = 1$

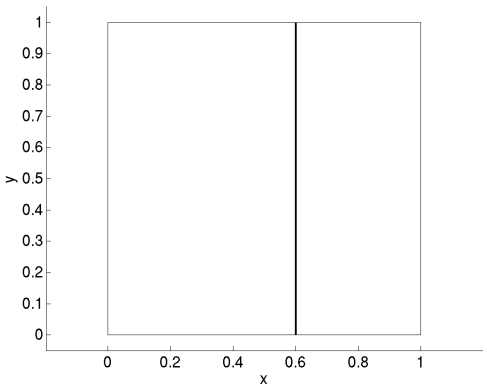

$t = 20$

Figure 12: Interface at specified times for $\lambda = \infty$.

4.2 Quantitative Accuracy and Efficiency Comparisons

The qualitative graphical results in the figures in the previous subsection indicate that the two component model appears to give results that are qualitatively consistent with the three species model in the asymptotic limit of $\lambda \to \infty$. This section defines four measures to study the accuracy of the two component model quantitatively. As our goal in this paper is to verify that the two component model is numerically superior to the three species model, we also define two measures to evaluate the efficiency quantitatively. The four measures of accuracy and two measures of efficiency are studied over a variety of finite element meshes with $N \times N$ elements.

The first measure of accuracy is the closeness of the location x^* of the interface at the final time of the transient problem to that of the actual steady state interface, which is $x^*_{\rm SS} = 0.601806640625$; this value was obtained by simulations of the three species steady state problem in 1-D on a high resolution mesh with $N = 8{,}192$ and will be considered the true value for x^* in the following. The steady state does not depend on y, hence the 1-D simulation of the steady state problem gives information on how our 2-D problem should behave at steady state. The other measures of accuracy are three times t_1, t_2, and t_3 defined by significant transitions in the development of the interface plot. The time t_1 represents the time at which the head separates from the body in the left portion of the domain; t_2 represents the time at which the head re-attaches to the body, and t_3 represents the time at which the disjoint disk in the upper right of the domain dissipates. These times were obtained by viewing movies of the interface, where the movies were compiled using the output of all time steps from the ODE solver, which are significantly more than the six times in the figures in the previous subsection. Thus the times listed are as precise as possible for each simulation, but are limited in accuracy by the number of time steps and their exact choice made by the time step selection algorithm in the ODE solver. The x^* values in Table 2 are identical for the $\lambda = 10^6$, 10^9, and ∞ cases. They are slightly different among the mesh resolutions, but within the mesh resolution to $x^*_{\rm SS}$ in each case. In $\lambda = 10^3$ case, the x^* values are slightly off, because they have not reached their asymptotic values, yet. We see in Table 2 that the corresponding times t_1, t_2, and t_3 are very close for all $\lambda = 10^3$, 10^6, 10^9, and ∞, for each fixed mesh resolution $N \times N$. The difference in the values of the times between different mesh resolutions is explained by the fact that they generate slightly different initial conditions, depending on how exactly the curved boundaries in the initial conditions are resolved, which result in slightly different behavior over time. Because the accuracy results are nearly identical across at least $\lambda = 10^6$, 10^9, and ∞ for each fixed mesh resolution, these accuracy measures confirm that the two component model ($\lambda = \infty$) is an accurate simulator for the three species model in the asymptotic limit.

Table 2: Summary of accuracy and efficiency data for simulations (a) of the three species model with $\lambda = 10^3$, (b) of the three species model with $\lambda = 10^6$, (c) of the three species model with $\lambda = 10^9$, and (d) of the two component model with $\lambda = \infty$.

(a) $\lambda = 10^3$						
$N \times N$	x^*	$t_1(\times 10^{-4})$	$t_2(\times 10^{-3})$	$t_3(\times 10^{-3})$	steps	time (s)
64×64	0.609375000	1.608304	7.077499	7.077499	215	43
128×128	0.609375000	4.961102	6.354030	6.768004	230	166
256×256	0.601562500	3.716557	6.972642	6.972642	249	771
(b) $\lambda = 10^6$						
$N \times N$	x^*	$t_1(\times 10^{-4})$	$t_2(\times 10^{-3})$	$t_3(\times 10^{-3})$	steps	time (s)
64×64	0.625000000	1.567014	6.863272	6.795620	573	140
128×128	0.609375000	4.831523	6.317239	6.768036	571	467
256×256	0.605468750	3.578299	6.753300	6.836752	548	2,463
(c) $\lambda = 10^9$						
$N \times N$	x^*	$t_1(\times 10^{-4})$	$t_2(\times 10^{-3})$	$t_3(\times 10^{-3})$	steps	time (s)
64×64	0.625000000	1.661473	6.862192	6.790343	646	231
128×128	0.609375000	4.804291	6.267175	6.696524	1,743	1,802
256×256	0.605468750	3.612976	6.685775	6.791933	2,353	11,391
(d) $\lambda = \infty$						
$N \times N$	x^*	$t_1(\times 10^{-4})$	$t_2(\times 10^{-3})$	$t_3(\times 10^{-3})$	steps	time (s)
64×64	0.625000000	1.564306	6.903623	6.903623	203	31
128×128	0.609375000	4.910450	6.459946	6.623560	229	95
256×256	0.605468750	3.619295	7.091791	7.091791	234	444

As measures of efficiency, we consider the number of time steps taken by the ODE solver and the total computation time in seconds (s) taken by the COMSOL. We see from Table 2 that the number of time steps for each λ value is on the same order of magnitude for all $N \times N$ meshes reported, except for the most numerically challenging $\lambda = 10^9$ case. The computation times get significantly larger for the finer meshes due to the larger linear systems that need to be solved in each time step. For each fixed N, as λ increases, the time of computation and the number of time steps increase rapidly for the finite λ values. But for two component model with $\lambda = \infty$ in Table 2, the number of time steps are again on the scale of the $\lambda = 10^3$ case. This indicates that the smoothness of the two component model is comparable to the three species model with this moderate λ value; the computation time is even faster than that case resulting from the smaller number of unknowns in the system for two PDEs in the two component model as compared to three for the three species model. We point out that in the case of the 64×64 mesh for $\lambda = 10^9$, the ODE solver failed to converge at the ODE tolerances stated above and used for all other cases. However, this was for a different reason than with the default ODE solver. The default solver BDF-IDA failed to find suitable convergence for the initial conditions, while in this case the ODE solver BDF-DASPK computed until it hit a point in time where it could not compute further given the minimum time step value. Thus, the data listed in the table for the 64×64 mesh for $\lambda = 10^9$ is computed with relative and absolute ODE tolerences that are one order of magnitude coarser than the ones used for all other cases.

Taken together, because the accuracy results are nearly identical across at least $\lambda = 10^6$, 10^9, and ∞ for each fixed mesh resolution, the accuracy and efficiency measures demonstrate that the two component model ($\lambda = \infty$) is an accurate and efficient simulator for the three species model in the asymptotic limit on a given mesh.

Acknowledgments

The work reported in this paper is based on consulting projects in the Center for Interdisciplinary Research and Consulting (`www.umbc.edu/circ`). The hardware used in the computational studies is part of the UMBC High Performance Computing Facility (HPCF). The facility is supported by the U.S. National Science Foundation through the MRI program (grant nos. CNS–0821258 and CNS–1228778) and the SCREMS program (grant no. DMS–0821311), with additional substantial support from the University of Maryland, Baltimore County (UMBC). See `www.umbc.edu/hpcf` for more information on HPCF and the projects using its resources.

References

[1] Aaron Churchill, Matthias K. Gobbert, and Thomas I. Seidman. Efficient computation for a reaction-diffusion system with a fast reaction in two spatial dimensions using COMSOL Multiphysics. Technical Report HPCF–2009–7, UMBC High Performance Computing Facility, University of Maryland, Baltimore County, 2009.

[2] Leonid V. Kalachev and Thomas I. Seidman. Singular perturbation analysis of a stationary diffusion/reaction system whose solution exhibits a corner-type behavior in the interior of the domain. *J. Math. Anal. Appl.*, vol. 288, pp. 722–743, 2003.

[3] Michael Muscedere and Matthias K. Gobbert. Parameter study of a reaction-diffusion system near the reactant coefficient asymptotic limit. *Dynamics of Continuous, Discrete and Impulsive Systems Series A Supplement*, pp. 29–36, 2009.

[4] W. Nernst. Theorie der Reaktionsgeschwindigkeit in heterogenen Systemen. *Z. Phys. Chem.*, vol. 47, pp. 52–55, 1904.

[5] Noemi Petra and Matthias K. Gobbert. Parallel performance studies for COMSOL Multiphysics using scripting and batch processing. In Yeswanth Rao, editor, *Proceedings of the COMSOL Conference 2009, Boston, MA*, 2009.

[6] Thomas I. Seidman. Interface conditions for a singular reaction-diffusion system. *Discrete and Cont. Dynamical Systems — Series S*, vol. 2, no. 3, pp. 631–643, 2009.

[7] Thomas I. Seidman and Leonid V. Kalachev. A one-dimensional reaction/diffusion system with a fast reaction. *J. Math. Anal. Appl.*, vol. 209, pp. 392–414, 1997.

[8] Thomas I. Seidman and Adrian Muntean. Fast-reaction asymptotics for a time-dependent reaction-diffusion system with a nonlinear source term. Technical Report CASA report 10-51, Technische Universiteit Eindhoven, 2010.

[9] Ana Maria Soane, Matthias K. Gobbert, and Thomas I. Seidman. Numerical exploration of a system of reaction-diffusion equations with internal and transient layers. *Nonlinear Anal.: Real World Appl.*, vol. 6, no. 5, pp. 914–934, 2005.