A STUDY OF THE NUMERICAL STABILITY OF AN IMEX SCHEME WITH APPLICATION TO THE POISSON-NERNST-PLANCK EQUATIONS

M. C. PUGH*, DAVID YAN , AND F. P. DAWSON†

Abstract. The Poisson-Nernst-Planck equations with generalized Frumkin-Butler-Volmer boundary conditions (PNP-FBV) describe ion transport with Faradaic reactions and have applications in a wide variety of fields. Using an adaptive time-stepper based on a second-order variable step-size, semi-implicit, backward differentiation formula (VSSBDF2), we observe that when the underlying dynamics is one that would have the solutions converge to a steady state solution, the adaptive time-stepper produces solutions that "nearly" converge to the steady state and that, simultaneously, the time-step sizes stabilize at a limiting size dt_∞. Linearizing the SBDF2 scheme about the steady state solution, we demonstrate that the linearized scheme is conditionally stable and that this is the cause of the adaptive time-stepper's behaviour. Mesh-refinement, as well as a study of the eigenvectors corresponding to the critical eigenvalues, demonstrate that the conditional stability is not due to a time-step constraint caused by high-frequency contributions. We study the stability domain of the linearized scheme and find that it can have corners as well as jump discontinuities.

Key words. Poisson-Nernst-Planck Equations, Semi-Implicit Methods, IMEX Methods, Adaptive time-stepping, Conditional Linear Stability

AMS subject classifications. 65M12, 65Z05, 78A57

1. Introduction. In this article, we study the numerical stability properties of a second-order implicit-explicit backwards differencing formula (SBDF2) as applied to the Poisson-Nernst-Planck equations with generalized Frumkin-Butler-Volmer boundary conditions (the PNP-FBV system). The PNP equations describe the transport of charged species subject to diffusion and electromigration. The Poisson-Nernst-Planck (PNP) equations describe the transport of charged species subject to diffusion and electromigration. They have wide applicability in electrochemistry, and have been used to model a number of different systems, including porous media [1, 2, 3, 4], micro-electrodes [5, 6], ion-exchange membranes [7, 8], electrokinetic phenomena [9, 10, 11], ionic liquids [12, 13], electrochemical thin films [14, 15, 16], fuel cells [17], supercapac-itors [18], and many more. The Frumkin-Butler-Volmer boundary conditions describe charge transfer reactions at electrodes.

The PNP-FBV system is a parabolic-elliptic system with nonlinear boundary conditions that model reactions at the electrodes. There is a singular perturbation parameter ϵ; small values of ϵ lead to thin boundary layers with sharp transitions to the behaviour in the bulk. In the parameter regimes of interest, the dynamics are strongly diffusive and so an implicit-explicit scheme is a natural approach: the linear diffusive term is handled implicitly and the nonlinear terms are handled explicitly.

If one considers a system of ODEs of the form $\mathbf{u}_t = \mathbf{L}(\mathbf{u}) + \mathbf{f}(\mathbf{u})$ and applies an r-step LMM to the system, this yields

$$\frac{1}{dt}\sum_{j=0}^{r} a_j \mathbf{u}_{n+j} = \sum_{j=0}^{r} c_j \mathbf{L}_{n+j} + b_j \mathbf{f}_{n+j} = \sum_{j=0}^{r} c_j \mathbf{A}\,\mathbf{u}_{n+j} + b_j \left(\mathbf{B}_{n+j} + \mathbf{f}_{n+j}\right)$$

where $\mathbf{u}_n \approx \mathbf{u}(n\,dt)$ approximates the true solution. The expression includes the option of an operator splitting: $\mathbf{L}(\mathbf{u}) = \mathbf{A}\mathbf{u} + \mathbf{B}(\mathbf{u})$, where $\mathbf{A}$ is a matrix. If $b_r =$

*Department of Mathematics, University of Toronto (mpugh@math.utoronto.ca)
†Department of Electrical and Computer Engineering, University of Toronto

$c_r = 0$, the scheme is explicit. Similarly, if $b_j = c_j = 0$ for $0 \leq j \leq r - 1$ the scheme is implicit. Of particular interest are semi-implicit schemes: $c_r \neq 0$, $b_r = 0$ and some c_j and/or b_j with $0 \leq j \leq r - 1$ are nonzero. These are used for nonlinear problems in which there are fast time-scales that yield numerical stability constraints on the time-steps if explicit schemes are used.

In a semi-implicit scheme, one solves for $\mathbf{u}_{n+r}$ at each time-step. If there is no splitting and $\mathbf{L}$ is nonlinear, this requires a (usually slow) nonlinear solve. If $\mathbf{L}$ is linear, or if an operator splitting is used, finding $\mathbf{u}_{n+r}$ requires solving a linear system. If the matrix does not have a "helpful" structure (such as being banded) or is a full matrix, this can require (possibly slow) iterative methods such as GMRES for the solution.

There have been many approaches to the challenge of "stiffness reduction" whether in semi-implicit linear multi-step methods [19, 20, 21, 22, 23, 24, 25], Runge-Kutta methods [26, 27, 28], matrix exponential/integrating factor methods [29, 30, 31], and other approaches [32, 33, 34, 35]. The cited articles are provided as seminal/well-written samples of a large literature on the topic.

The article is not about stiffness reduction. Rather, it's about a system in which the stiffness is well-handled by a commonly-used ImEx scheme. However, as the solution equilibrates, the scheme becomes conditionally stable but not due to fast time-scales at high frequencies. It is likely that some of the methods proposed in the cited works would allow one to use knowledge of the structure of the linearized operator about the steady-state solution so as to modify the scheme and shift it out of the conditional stability.

For the PNP-FBV system, in [36, 37, 38], we developed and used an adaptive time-stepping scheme based on a second-order variable step-size, semi-implicit, backward differentiation formula (VSSBDF2 [39]). We considered a variety of applied voltages and currents; the adaptive time-stepper was vital in that it could refine, and subsequently coarsen, the time-steps in response to fast changes in the applied forcing. However, when the applied forcing was held constant, and the underlying physical solution relaxes to a steady state solution, the adaptive time-stepper did not behave in the expected manner: the time-steps didn't grow until they reached the user-specified maximum time-step size, $dt_{\max}$, and the numerical solution didn't converge to the numerical steady state. Instead, we observe that the numerical solution gets close to, but fails to converge to, the numerical steady state and, simultaneously, the time-step sizes stabilize at a limiting step size dt_∞. In this regime, the VSSBDF2 adaptive time-stepper is effectively an SBDF2 time-stepper with time-steps equal to dt_∞.

Rosam, Jimack and Mullis [40] used an adaptive SBDF2 algorithm to study a problem in binary alloy solidification. In their Figure 4, they appear to show time-steps stabilizing to a constant value (i.e. thresholding), but the reason is not given: they report that it is related to the tolerance set in the adaptive time-stepper. We did not find such an phenomenon when we varied the tolerance; we found the same limiting time-step size dt_∞.

In this work, we perform a stability analysis of the SBDF2 scheme linearized about the steady state solution. We demonstrate that the linearized scheme[1] is conditionally

---

[1] There are articles in which an ODE or PDE is rewritten in a way that makes a certain numerical schemes perform better. For example, $u_t = u\, u_x$ would be "linearized" by being written as $u_t = \alpha\, u_x + (u - \alpha)\, u_x$, the time-stepping scheme would treat the $\alpha\, u_x$ term implicitly and the $(u - \alpha)\, u_x$

stable with a stability threshold dt^*. We present simulations demonstrating that $dt_\infty = dt^*$; this suggests the VSSBDF2 adaptive time-stepper is finding the stability threshold. Depending on the physical parameter values, when $dt = dt^*$ there's one one eigenvector with eigenvalue -1 or there's a pair of eigenvectors with complex eigenvalues of magnitude 1. The eigenvectors are not highly oscillatory and when performing the stability analysis using different spatial discretizations we find that the stability threshold, dt^*, does not change significantly. Specifically, dt^* doesn't go to zero as dx goes to zero; this shows that the conditional stability is not of "diffusive type" in which high frequencies can grow exponentially in time if dt is too large.

The observed conditional stability is problem-specific. We used the VSSBDF2 adaptive time-steppers on several dissipative systems that had non-trivial steady states and found that simple systems such as reaction diffusion equations did not have asymptotically stable steady states which yielded conditionally stable schemes when the SBDF2 scheme was linearized about them. We did find that toy models of the PNP-BDF system did have steady states which yielded conditionally stable schemes when the SBDF2 scheme was linearized about them [37].

The PNP-FBV system has a singular perturbation parameter ϵ. By varying ϵ, we are able to study the stability domain to see how dt^* depends on ϵ. We find that the stability domain is not smooth — there can be corners and jump discontinuities in the graph of $dt^*(\epsilon)$. Jump discontinuities are especially striking because they mean that the same value of dt could yield a stable SBDF2 computation for one value of ϵ but could result in a computation that blows up for another, close value of ϵ. We have not seen this type of behaviour reported in the literature. We find that for small values of ϵ the stability domain is not significantly influenced by the value of the (constant) applied voltage or current.

The structure of the stability domain is also problem-specific. We considered some other dissipative systems and didn't find stability domains with corners, cusps, or jumps; otherwise we would present results for a simpler system than the PNP-FBV system.

The VSSBDF2 adaptive time-stepper can be used with, or without, a Richardson extrapolation step. Richardson extrapolation is a common way to increase the accuracy of a scheme. We demonstrate that Richardson extrapolation can affect the linear stability of a scheme in various ways. We give an example of a scheme that is unconditionally stable when there is no Richardson extrapolation step and is conditionally stable when there is one.

Our methods are not restricted to the PNP-FBV system or to the VSSBDF2 adaptive time-stepper. If one is using a linear multistep method (LMM) to study a physical system that has asymptotically stable steady states, our approach is relevant in the following ways. 1) We give a heuristic argument based on the local truncation error as to why, in general, an adaptive time-stepper would naturally find the stability threshold if the underlying constant-time-step scheme is conditionally stable when linearized about the steady state. As a result, if one builds an adaptive time-stepper based on a variable-step-size version of the LMM being used, one can use the adaptive time-stepper to explore the stability domain by computing the initial value problem. If the time-step size stabilizes at a value dt_∞, this suggests that the underlying constant-time-step scheme is conditionally stable with stability threshold

---

term explicitly, and by choosing α carefully, good things happen. This is not the type of linearization we're referring to.

dt^* and $dt_\infty = dt^*$. 2) If the adaptive time-stepper reveals that the constant-time-step scheme is conditionally stable when linearized about the steady state, and one wishes to compute the numerical coarse steady state solution up to round-off error, one can do this with confidence by using the constant-time-step scheme with a time-step size chosen smaller than dt_∞. 3) Although the adaptive time-stepper can be used to explore the stability domain of the constant-time-step scheme when linearized about the steady state, one still needs to linearize the scheme about the steady state and study the eigenvalues of the linearized problem in order to to understand stability domain features such as corners, jumps, and whether or not it's a single real-valued eigenvalue that goes unstable as dt exceeds dt^*. The procedure we use to linearize about the steady state and find the eigenvalues and eigenvectors of the linearized system could be used for any LMM.

2. The adaptive time-stepper. Consider the ODE $u' = f(u) + g(u)$ where $f(u)$ is a nonlinear term and $g(u)$ is a stiff linear term. Given u^{n-1} at time $t^{n-1} = t^n - dt_{\text{old}}$ and u^n at time t^n, u^{n+1} at time $t^{n+1} = t^n + dt_{\text{now}}$ is determined via

(1) SBDF2: $\quad \dfrac{1}{dt}\left(\dfrac{3}{2}u^{n+1} - 2\,u^n + \dfrac{1}{2}u^{n-1}\right) = 2\,f(u^n) - f(u^{n-1}) + g(u^{n+1}),$

where the superscript notation denotes time levels: u^n approximates $u(t^n)$ (see, for example, [41]). Our VSSBDF2 adaptive time-stepper is based on a second-order variable step-size implicit-explicit backwards differencing formula, introduced by Wang and Ruuth [39], as a generalization of the SBDF2 scheme:

VSSBDF2: $\quad \dfrac{1}{dt_{\text{now}}}\left(\dfrac{1+2\omega}{1+\omega}u^{n+1} - (1+\omega)u^n + \dfrac{\omega^2}{1+\omega}u^{n-1}\right)$

(2) $$= (1+\omega)f(u^n) - \omega f(u^{n-1}) + g(u^{n+1}),$$

where $\omega = dt_{\text{now}}/dt_{\text{old}}$.

The VSSBDF2 adaptive time-stepper is described in detail in the companion article [37, 38]. The key idea is: if one has computed the (approximate) solution up to the current time, (u^l, t^l) for $l = 0, \ldots, n$, one can use these solutions and the time-stepper to choose a new time t^{n+1} so that the local truncation error $\|u^{n+1} - u(t^{n+1})\|$ is "small but not too small".

One can't know the local truncation error if one doesn't know the (exact) solution $u(t^{n+1})$; in practice one needs an approximation of the local trucation error. We do this as follows. First, we choose a candidate time-step: $dt_{\text{now}} = dt_{\text{old}}$, for example. We then take one "coarse" step from $t^n - dt_{\text{old}}$ and t^n to $t^n + dt_{\text{now}}$, using u^{n-1} and u^n to create u_c^{n+1}. Next, we take one "fine" step from $t^n - dt_{\text{old}}/2$ and t^n to $t^n + dt_{\text{now}}/2$, using $u_f^{n-1/2}$ and u^n to create $u_f^{n+1/2}$ and take a second fine step from t^n and $t^n + dt_{\text{now}}/2$, using u^n and $u_f^{n+1/2}$ to create u_f^{n+1}. We use u_c^{n+1} and u_f^{n+1} to estimate the local truncation error [37, 38]:

(3) $$\epsilon_c^{n+1} = \dfrac{8\,(dt_{\text{old}} + dt_{\text{now}})}{7dt_{\text{old}} + 5dt_{\text{now}}}\left(u_c^{n+1} - u_f^{n+1}\right) \approx u_c^{n+1} - u(t^{n+1}).$$

If the error is acceptable we advance in time. If the error is unacceptable we choose a new dt_{now} and try again. The time-steps are constrained to by the user-specified lower and upper bounds: dt_{min} and dt_{max}. If dt_{now} has been accepted, we construct

u^{n+1} in one of two ways:

(4)   No Richardson Extrapolation: $u^{n+1} = u_c^{n+1}$

(5)   With Richardson Extrapolation: $u^{n+1} = \alpha\, u_c^{n+1} + \beta\, u_f^{n+1}$

where

(6) $$\alpha = -\frac{dt_{\text{old}} + 3\, dt_{\text{now}}}{7\, dt_{\text{old}} + 5\, dt_{\text{now}}}, \quad \beta = 8\, \frac{dt_{\text{old}} + dt_{\text{now}}}{7\, dt_{\text{old}} + 5\, dt_{\text{now}}}.$$

The local truncation error for u^{n+1} defined by (5) is one order higher than the local truncation error for u^{n+1} defined by (4) [37, 38]. Note that if $dt_{\text{now}} = dt_{\text{old}}$ then (5) reduces to the standard Richardson extrapolation formula for second-order schemes.

2.1. The Logistic Equation. As a simple example to show how the VSSBDF2 adaptive time-stepper finds the stability threshold of the SBDF2 scheme, we consider the logistic equation

(7) $$u_t = r\, u\, (1 - u) = r\, u - r\, u^2$$

If $r > 0$ then the $u(t) = 0$ steady state is unstable and the $u(t) = 1$ steady state is asymptotically stable. Applying SBDF2 to (7) with $f(u) = -r\, u^2$ and $g(u) = r\, u$ yields the scheme

(8) $$\frac{1}{dt}\left(\frac{3}{2}u^{n+1} - 2u^n + \frac{1}{2}u^{n-1}\right) = ru^{n+1} - 2r(u^n)^2 + r(u^{n-1})^2.$$

Time-stepping this scheme using SBDF2 with relatively small time-steps, we find that solutions with positive initial data converge exponentially fast to the $u(t) = 1$ steady state. However, if the time-steps are too large, solutions fail to converge. See the left plot in Figure 1. If we use the VSSBDF2 adaptive time-stepper, we find that the solution comes close to $u(t) = 1$ but does not converge to it. The top middle plot in Figure 1 presents $\log(|1 - u(t)|)$ as a function of time. After a short transient, the solution decays roughly exponentially to $u(t) = 1$ but around time $t = 50$ the solution stays about 5×10^{-7} away from the steady state. The bottom middle plot presents the time-step size, dt as a function of time; we see that, after a transient, the time-step coarsens exponentially fast. It then refines and stabilizes at a time-step size of approximately $dt_\infty = 4/(7r)$.

To understand the time-step stabilization found in the VSSBDF2 adaptive time-stepper and the non-convergence found when using "large" time steps in the SBDF2 IMEX scheme, we linearized the SBDF2 scheme (8) about the constant steady state $u^n = 1$:

(9) $$\frac{1}{dt}\left(\frac{3}{2}d^{n+1} - 2d^n + \frac{1}{2}d^{n-1}\right) = r\, d^{n+1} - 4\, r\, d^n + 2\, r\, d^{n-1}.$$

above, d^l denotes the deviation from the steady state: $u^l = d^l + 1$. The linear stability analysis of Appendix A, implies (9) is conditionally stable with a stability threshold of $dt^* = 4/(7r)$.

One demonstration that the stability threshold, dt^*, of the linearized scheme (9) is relevant to the SBDF2 time-stepper (8) is that the VSSBDF2 adaptive time-stepper had its time-steps stabilize to dt^* as it tried, but failed, to reach the steady state
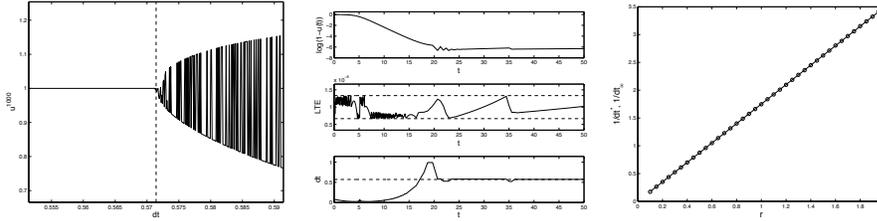
Fig. 1: The logistic equation (7) is simulated with $r = 1$ and initial data $u(0) = .01$. Left Figure: For each (fixed) value of dt, we time-step using SBDF2 until the 1000th iterate has been found. (The first step is taken using Forward Euler.) In the figure, u^{1000} is plotted against dt. The dashed line indicates the stability threshold $dt^* = 4/(7r)$. Middle figure: The VSSBDF2 adaptive time-stepper is used with $tol = 10^{-6}$ and $range = tol/3$. The stability constraint is $dt^* = 4/(7r) \approx 1.14$; $dt_{\max} = 2dt^*$. Top plot: $\log(1 - u(t))$ versus time. Middle plot: local truncation error as a function of time. Dashed lines indicate the interval $(tol - range, tol + range)$. Bottom plot: time-step size versus time. Dashed line indicates $dt^* = 4/(7r)$. Right figure: The VSSBDF2 adaptive time-stepper for a range of values of r. The stability constraint is $dt^* = 4/(7r)$; we take $dt_{\max} = 2dt^*$. For each value of r, the solution is computed to $t = 750$ and dt_∞ is found by averaging the final 100 values of dt. In the figure, $1/dt_\infty$ is plotted versus r with circles and $1/dt^*$ is plotted with a solid line.

solution $u(t) = 1$. See the bottom middle plot figure in Figure 1. As a second demonstration, the left plot of Figure 1 demonstrates that if $dt < dt^*$ then the SBDF2 time-stepper (8) yields solutions that converge to the steady state $u^n = 1$ while if $dt > dt^*$ solutions do not. This is unsurprising because the linearized scheme (9) will determine the dynamics for any iterate that is close to the steady state $u^n = 1$. Finally, we varied the parameter r and studied how the limiting time-step size found by the VSSBDF2 adaptive time-stepper dt_∞ is related to the stability threshold, dt^*, of the linearized scheme (9). The right plot of Figure 1 demonstrates that the VSSBDF2 adaptive time-stepper is sharply finding the stability threshold time-step size.

3. Simulations of the PNP-FBV system. The logistic equation provided a simple example where the SBDF2 method, when linearized about the asymptotically stable steady state, is conditionally stable. As demonstrated, the VSSBDF2 adaptive time-stepper finds the stability threshold. The stability domain for the logistic equation is simple: $dt^* = 4/(7r)$; this is a smooth curve when plotted as a function of r.

We now turn to the PNP-FBV system. The stability domain of the SBDF2 method, linearized about the asymptotically stable steady state, is neither smooth nor continuous; see Figure 5. The PNP-FBV system is significantly more complex than the logistic equation and a natural question whether there are simpler dissipative systems for which the SBDF2 method had stability domains with corners, cusps, or jumps. We tried a few simple options such as reaction diffusion equations and simple models of the PNP-FBV system with different types of boundary conditions [37] but the stability domains were not as interesting.

The one-dimensional, nondimensionalized PNP equations for a media with 2 mo-

bile species is

$$(10) \qquad \frac{\partial c_\pm}{\partial t} = -\frac{\partial}{\partial x}\left[-\frac{\partial c_\pm}{\partial x} - z_\pm\, c_\pm\, \frac{\partial \phi}{\partial x}\right], \qquad t > 0,\ x \in (0,1),$$

$$(11) \qquad -\epsilon^2 \frac{\partial^2 \phi}{\partial x^2} = \frac{1}{2}\left(z_+ c_+ + z_-\, c_-\right), \qquad\qquad x \in (0,1),$$

where $c_\pm$ and $z_\pm$ are the concentration and charge number of the positive/negative ion, ϕ is the potential and ϵ is the ratio of the Debye screening length to the interelectrode width L. The PNP equations are combined with generalized Frumkin-Butler-Volmer (FBV) boundary conditions[2] on each species which describe charge transfer reactions at electrodes. We choose $z_+ = -z_- = 1$ with two no-flux boundary conditions and two FBV boundary conditions for equation (10) :

$$(12)$$
$$-\left(-\frac{\partial c_-}{\partial x} - c_-\,\frac{\partial \phi}{\partial x}\right)\bigg|_{x=0} = \left(-\frac{\partial c_-}{\partial x} - c_-\,\frac{\partial \phi}{\partial x}\right)\bigg|_{x=1} = 0$$

$$(13)$$
$$-\left(-\frac{\partial c_+}{\partial x} - c_+\,\frac{\partial \phi}{\partial x}\right)\bigg|_{x=0} = F(t) := 4k_{c,a}\,c_+(0,t)\,e^{-0.5\,\Delta\phi_{\text{left}}} - 4\,j_{r,a}\,e^{0.5\,\Delta\phi_{\text{left}}},$$

$$(14) \quad \left(-\frac{\partial c_+}{\partial x} - c_+\,\frac{\partial \phi}{\partial x}\right)\bigg|_{x=1} = G(t) := 4k_{c,c}\,c_+(1,t)\,e^{-0.5\,\Delta\phi_{\text{right}}} - 4\,j_{r,c}\,e^{0.5\,\Delta\phi_{\text{right}}},$$

where $k_{c,a}$, $k_{c,c}$, $j_{r,a}$, and $j_{r,c}$ are reaction rate parameters; the second part of the subscripts a and c refer to the anode and cathode, respectively . Equations (13)–(14) model the electrodeposition reaction $\text{C}^+ + \text{e}^- \rightleftharpoons \text{M}$ where M represents the electrode material. There is a compact layer of charge, called the Stern layer, that occurs in the electrolyte next to an electrode surface [42, 43]. In equations (13)–(14), $\Delta\phi_{\text{left}}$ and $\Delta\phi_{\text{right}}$ refer to the potential differences across the Stern layers that occur at the anode and cathode respectively. Specifically,

$$(15)\quad \Delta\phi_{\text{left}} = \phi_{\text{anode}} - \phi(0,t) = -\phi(0,t),\ \Delta\phi_{\text{right}} = \phi_{\text{cathode}} - \phi(1,t) = v(t) - \phi(1,t)$$

where the potential at the anode has been set to zero and $v(t)$ denotes the potential at the cathode. In addition, the Poisson equation (11) uses a mixed (or Robin) boundary condition [14, 15, 16],

$$(16) \qquad -\epsilon\,\delta\,\frac{\partial \phi}{\partial x}\bigg|_{x=0} = \Delta\phi_{\text{left}} := -\phi(0,t),$$

$$(17) \qquad +\epsilon\,\delta\,\frac{\partial \phi}{\partial x}\bigg|_{x=1} = \Delta\phi_{\text{right}} := v(t) - \phi(1,t),$$

where δ is the ratio of the compact layer thickness to L. Finally, there is an ODE which ensures conservation of electrical current at the electrode [44, 45],

$$(18) \qquad -\frac{\epsilon^2}{2}\frac{d}{dt}\phi_x(1,t) = j_{\text{ext}}(t) - \left[k_{c,c}\,c_+(1,t)\,e^{-0.5\,\Delta\phi_{\text{right}}} - j_{r,c}\,e^{0.5\,\Delta\phi_{\text{right}}}\right],$$

where $j_{\text{ext}}(t)$ is the current through the device. We refer to the PNP equations with the generalized Frumkin-Butler-Volmer boundary conditions as the PNP-FBV systsem.

---

[2]written here, for example, for a reaction involving one ionic species and one electron

The device is operated in two regimes — either the current or the voltage at the cathode is externally controlled. If the voltage at the cathode, $v(t)$, is externally controlled then the the PNP-FBV system (10)–(11) with boundary conditions (12)–(14) and (15)–(17) are numerically solved, determining $c_\pm$ and ϕ. The current is found a postiori using equation (18). If the current, $j_{\text{ext}}(t)$, is externally controlled then equation (18) is part of the PNP-FBV system and the ODE is numerically solved along with the PDEs, determining $c_\pm$, ϕ, and $\phi_x(1,t)$ simultaneously. The voltage $v(t)$ is then found a postiori.

The companion article [38] presents the numerical scheme in full: spatial discretization, boundary conditions, splitting scheme, and error control. For reasons discussed at the end of Section 4, unless noted otherwise, the simulations presented in the rest of this section and in Section 4 did not use Richardson extrapolation in the VSSBDF2 adaptive time-stepper. Specifically, u^{n+1} was defined using (4).

Figure 2 considers an initial value problem for the PNP-FBV system (10)–(17) with constant imposed voltage. The initial data is fixed, as are all the other physical parameters. Solutions are computed using the VSSBDF2 adaptive time-stepper.

The top figure demonstrates that, after a short transient, the solution initially decays exponentially quickly to a numerical steady state. However, once the solution is within (approximately) 10^{-7} of the steady state solution, this convergence ends and the computed solution stays about 10^{-7} away from the steady state. The middle figure demonstrates that the VSSBDF2 adaptive time-stepper is keeping the (approximate) local truncation error (3) within the user-specified range of $(tol - range, tol + range)$. The bottom figure demonstrates that the time-step size initially increases exponentially fast and after a while it decreases and stabilizes to dt_∞. The dashed line in the bottom figure is the stability threshold found by the linear stability analysis discussed in Section 4: dt^*. This simulation demonstrates that the VSSBDF2 adaptive time-stepper appears to eventually stabilize at a time-step size that is precisely the stability threshold.

The top left plot presents the deviation from the numerical steady state $\mathbf{c}_{+,ss}$. The numerical steady state, $\mathbf{c}_{\pm,ss}$ and $\boldsymbol{\phi}_{ss}$, satisfies the discretized version of $0 = c_{\pm,xx} + z_\pm (c_\pm \phi_x)$ and (11). The numerical steady state was found by first computing an initial value problem using the VSSBDF2 adaptive time-stepper and then repeating the computation using the SBDF2 time-stepper with a (fixed) time-step size that is less than the value dt_∞ found by the adaptive time-stepper. We find that a long time simulation using fixed time-steps results in a solution that relaxes to a steady state solution as long as the time-steps are taken to be smaller than dt_∞. We use the solution at a late time as the numerical steady state solution.

4. Numerical Linear Stability. As demonstrated in Figure 2, if the applied voltage or applied current is constant for a period of time and if the user-specified maximum time-step $dt_{\max}$ is sufficiently large then the VSSBDF2 adaptive time-stepper stabilizes to take (nearly) constant time steps dt_∞ with $dt_\infty < dt_{\max}$. To understand this phenomenon, we study the linear stability of the SBDF2 time-stepper; this will give us insight into what the VSSBDF2 adaptive time-stepper is doing when it is taking (nearly) constant time-steps.

Consider the SBDF2 scheme (1) applied to the system of ODEs that arises from spatially discretizing the PDE $u_t = f(u, u_x, u_{xx}, \dots) + g(u, u_x, u_{xx}, \dots)$

$$(19) \qquad \frac{1}{dt}\left(\frac{3}{2}\mathbf{u}^{n+1} - 2\mathbf{u}^n + \frac{1}{2}\mathbf{u}^{n-1}\right) = \mathbf{g}(\mathbf{u}^{n+1}) + 2\mathbf{f}(\mathbf{u}^n) - \mathbf{f}(\mathbf{u}^{n-1}).$$
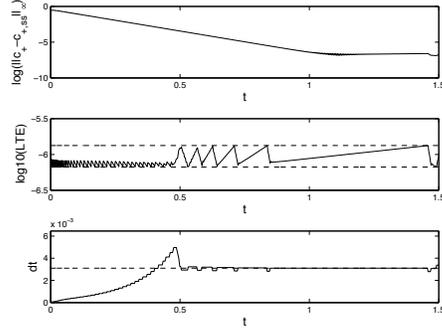
Fig. 2: PNP-FBV system (10)–(17) with constant voltage $v(t) = 2$, $\epsilon = .05$, and all other physical parameters set to 1. The initial data is $c_\pm(x, 0) = 1 + .1\sin(2\pi x)$ and $\phi_x(1, 0) = 0$. The simulation parameters (see [38]) are $N = 90$, $L_1 = L_2 = L_3 = 1/3$ and the adaptive time-stepping parameters are with $tol = 10^{-6}$, $range = tol/3$, and $dt_{\max} = 1$. The VSSBDF2 adaptive time-stepper is used. Top plot: Comparison of the solutions found by the VSSBDF2 adaptive time-stepper to the pre-computed numerical steady-state solution. Plot is $\log(\|\mathbf{c}_+^n - \mathbf{c}_{+,ss}\|_\infty)$ versus t_n. Deviations of $\mathbf{c}_-$ and ϕ from the corresponding steady state profiles behave similarly. Middle plot: The logarithm of the approximate local truncation error, (3), is plotted versus time. The dashed lines indicate $\log(tol \pm range)$. Bottom plot: Time-step size, dt, plotted versus time. The dashed line indicates the stability threshold $dt^* = 3.1000 \times 10^{-3}$ computed using the linear stability analysis presented in Section 4.

A steady state satisfies $0 = \mathbf{g}(\mathbf{u}_{ss}) + \mathbf{f}(\mathbf{u}_{ss})$. Linearizing about $\mathbf{u}_{ss}$ yields

$$(20) \quad \frac{1}{dt}\left(\frac{3}{2}\mathbf{d}^{n+1} - 2\mathbf{d}^n + \frac{1}{2}\mathbf{d}^{n-1}\right) = \mathbf{J_g}(\mathbf{u}_{ss})\,\mathbf{d}^{n+1} + 2\,\mathbf{J_f}(\mathbf{u}_{ss})\,\mathbf{d}^n - \mathbf{J_f}(\mathbf{u}_{ss})\,\mathbf{d}^{n-1}.$$

where $\mathbf{d}^m = \mathbf{u}^m - \mathbf{u}_{ss}$ are the deviations from the steady state and $\mathbf{J_f}(\mathbf{u}_{ss})$ and $\mathbf{J_g}(\mathbf{u}_{ss})$ are the Jacobian matrices evaluated at $\mathbf{u}_{ss}$; e.g. $(\mathbf{J_f}(\mathbf{u}_{ss}))_{ij} = \frac{\partial f_i}{\partial u_j}(\mathbf{u}_{ss})$. For simple problems, $\mathbf{J_f}$ and $\mathbf{J_g}$ can be determined analytically and evaluated at $\mathbf{u}_{ss}$. Otherwise, one can numerically approximate $\mathbf{J_f}(\mathbf{u}_{ss})$ and $\mathbf{J_g}(\mathbf{u}_{ss})$ in a variety of ways. We used a simple centre difference scheme. For example, the first column of $\mathbf{J_f}(\mathbf{u}_{ss})$ is approximated by

$$\frac{\mathbf{f}(\mathbf{u}_{ss} + h\,\mathbf{e}_1) - \mathbf{f}(\mathbf{u}_{ss} - h\,\mathbf{e}_1)}{2h}$$

where $\mathbf{e}_1$ is the first standard basis vector in $\mathbb{R}^N$: $(\mathbf{e}_1)_j = \delta_{j1}$ and h is relatively small.

For the PNP-FBV system (10)–(17), we find that the Jacobians cannot be simultaneously diagonalized. As a result, we cannot proceed based on an understanding of the time-stepper's behaviour for the ODE $y' = \mu\, y$. For this reason, we proceed with a direct computation of the eigenvalues and eigenvectors of the linearized scheme (20). One could do significantly more analysis of the stability if the Jacobians could be simultaneously diagonalized.

We rewrite the scheme (20) as

$$(21) \qquad\qquad \mathbf{d}^{n+1} = M_{\text{new}}\, M_{\text{now}}\, \mathbf{d}^n + M_{\text{new}}\, M_{\text{old}}\, \mathbf{d}^{n-1}$$

where

$$M_{\text{new}} = \left(\frac{3}{2}\,I - dt\,\mathbf{J_g}(\mathbf{u}_{ss})\right)^{-1}, M_{\text{now}} = 2\,I + 2\,dt\,\mathbf{J_f}(\mathbf{u}_{ss}), M_{\text{old}} = -\frac{1}{2}\,I - dt\,\mathbf{J_f}(\mathbf{u}_{ss}).$$

(21) is a system of N second-order linear difference equations. Solving it requires the initial deviation, $\mathbf{d}^0$, as well as the deviation after one time-step, $\mathbf{d}^1$. The system is rewritten [46, §D.2.1] as $2N$ first-order linear difference equations in the standard manner: $\mathbf{D}^n := [\mathbf{d}^{n-1}; \mathbf{d}^n] \in \mathbb{R}^{2N}$ and A is the companion matrix for the difference equation:

$$(22) \qquad \mathbf{D}^{n+1} = A\,\mathbf{D}^n = \begin{pmatrix} 0 & I \\ M_{\text{new}}\,M_{\text{old}} & M_{\text{new}}\,M_{\text{now}} \end{pmatrix} \mathbf{D}^n.$$

If $(\lambda_j, \mathbf{v}_j)$ is an eigenvalue-eigenvector pair of A then the structure of A implies that $\mathbf{v}_j = [\mathbf{d}_j; \lambda_j\,\mathbf{d}_j]$ for some $\mathbf{d}_j \in \mathbb{R}^N$. If A has $2N$ linearly independent eigenvectors, it follows that the general solution of the linearized problem (20) is

$$(23) \qquad \mathbf{d}^n = \sum_{j=1}^{2N} c_j\,\lambda_j^n\,\mathbf{d}_j$$

where the $2N$ coefficients, c_j, are determined using $\mathbf{d}^0, \mathbf{d}^1 \in \mathbb{R}^N$.

The connection between the linearized scheme (20) and the nonlinear scheme (19) is via the stability theory of fixed points for discrete dynamical systems. If $\mathbf{J_f}$ and $\mathbf{J_g}$ are continuous in a neighbourhood of $\mathbf{u}_{ss}$ and if dt is such that $(\frac{3}{2}\,I - dt\,\mathbf{J_g}(\mathbf{u}_{ss}))$ is invertible then the discrete dynamical system

$$(24) \quad \mathbf{U}^{n+1} = \mathbf{F}(\mathbf{U}^n) := \begin{pmatrix} \mathbf{U}_2^n \\ \left(\frac{3}{2}\,I - dt\,\mathbf{g}\right)^{-1}\left(2\mathbf{U}_2^n + 2\,dt\,\mathbf{f}(\mathbf{U}_2^n) - \frac{1}{2}\mathbf{U}_1^n - dt\,\mathbf{f}(\mathbf{U}_1^n)\right) \end{pmatrix}$$

is defined in a neighbourhood of the fixed point $[\mathbf{u}_{ss}\ \mathbf{u}_{ss}]$. Defining $\mathbf{U}^{n+1} = [\mathbf{u}^n; \mathbf{u}^{n+1}]$, this discrete dynamical system (24) is equivalent to the SBDF2 time-stepping scheme (19). The companion matrix A is the linearization of (24) at the fixed point $[\mathbf{u}_{ss}\ \mathbf{u}_{ss}]$. Therefore, if all eigenvalues of A have magnitude less than 1, then $[\mathbf{u}_{ss}\ \mathbf{u}_{ss}]$ is an asymptotically stable fixed point of (24) and $\mathbf{u}_{ss}$ is an asymptotically stable fixed point of the SBDF2 scheme (19).

4.1. Finding the stability threshold dt^* (if there is one). First, we need to find the (asymptotically stable) numerical steady state $\mathbf{u}_{ss}$. We do this by computing the initial value problem and finding the long-time limit.

Given a particular PDE (or set of PDEs) and boundary conditions, if the (constant time-step size) SBDF2 time-stepper (19) suggests that solutions of the initial value problem converge to a steady-state solution(s), one takes the long-time limit of the initial value problem as the steady-state solution $\mathbf{u}_{ss}$. One then computes (or approximates) the Jacobians about the steady state and constructs the matrix A using the time-step dt that was used in the SBDF2 time-stepper to find $\mathbf{u}_{ss}$. The eigenvalues of A will have magnitude less than one — otherwise the time-stepper wouldn't have found the steady-state solution. To determine if there is a linear stability constraint, one increases the time-step size, recomputes A and its eigenvalues, and then determines if any eigenvalues have magnitude greater than one for this new value of dt. Proceeding in this way, one seeks a time-step size for which an eigenvalue(s) crosses from magnitude less than one to magnitude greater than one. We then approximate the critical time-step size, dt^*, using a bisection method.

In practice, we don't have a priori knowledge of what would be a sufficiently small value of dt to use in the SBDF2 time-stepper to seek a steady-state solution. For this reason, we first use the VSSBDF2 adaptive time-stepper and if it yields a solution

that appears to be "trying" to converge to a steady state and the time-step sizes are stabilizing while keeping the local truncation error in the user-specified interval, we take this as evidence of an asymptotically stable steady state. We then repeat the simulation with the SBDF2 time-stepper, using a (constant) time-step that is smaller than the threshold value found by the VSSBDF2 adaptive time-stepper, to seek a steady state and proceed with the above stability study. If the time-steps increase to the user-specified $dt_{\max}$, one can repeat the simulation with larger and larger values of $dt_{\max}$ to see if the adaptive time-stepper stabilizes its timesteps at a value smaller than $dt_{\max}$. Or one can compute the numerical steady state $\mathbf{u}_{ss}$ and perform the linear stability analysis with larger and larger values of dt in the matrix A to seek a stability threshold. But if doing either of these does not identify a stability threshold, this isn't a proof that the SBDF2 scheme is unconditionally stable near the steady state; it's simply suggestive evidence.

Turning to the PNP-FBV system (10)–(17), for a fixed voltage and ϵ, we use the (constant time-step) SBDF2 time-stepper to find the steady state solution $\mathbf{c}_{+,ss}$, $\mathbf{c}_{-,ss}$, and $\boldsymbol{\phi}_{ss}$. The steady states of the evolution equation (10) are concatenated into one steady state $\mathbf{u}_{ss} := [\mathbf{c}_{+,ss}, \mathbf{c}_{-,ss}]$. The right-hand sides of the evolution equations are similarly concatenated: $\mathbf{f}$ is the spatial discretization of $[(c_+ \phi_x)_x, (c_- \phi_x)_x]$ and $\mathbf{g}$ is the discretization of $[c_{+,xx}, c_{-,xx}]$. We then approximate the Jacobians $\mathbf{J_f}(\mathbf{u}_{ss})$ and $\mathbf{J_g}(\mathbf{u}_{ss})$. If there are N mesh points then $\mathbf{u}_{ss} \in \mathbb{R}^{2N}$ and the Jacobians are $2N \times 2N$ matrices. A value of dt is chosen and the $4N \times 4N$ matrix A in (22) is constructed and its eigenvalues and eigenvectors are computed. The value of dt is then increased and the process is repeated.

Figure 3 presents results for a small value of ϵ with constant applied voltage. In the left figure, the magnitudes of all eigenvalues are plotted — we see that for small values of dt, all eigenvalues have magnitude less than one and that, as dt is increased, one branch goes unstable. We follow this branch to find the time-step size at which the eigenvalue's magnitude first equals 1: the stability threshold dt^*. For the ϵ of Figure 3, we find that one eigenvalue crosses the unit circle, crossing at value -1. In the top plot in the right figure, we plot the steady states $\mathbf{c}_{\pm,ss}$ and in the bottom plot in the right figure we plot the eigenfunctions at the stability threshold dt^*.

Figure 2 presents the results of applying the VSSBDF2 adaptive time-stepper to an initial value problem with $\epsilon = .05$. To demonstrate that the "barely stable" eigenvalue-eigenvector pair is the cause of the failure to converge to the steady state shown in the top plot of Figure 2, we took the computed solutions for $\mathbf{c}_+$ and $\mathbf{c}_-$ at a late time ($t = 100$) and computed the corresponding deviations from the steady state $\mathbf{d}_+$ and $\mathbf{d}_-$. These are presented in the top-right plot of Figure 3. In the bottom-right plot of Figure 3, the normalized deviations are plotted with open circles — note that they closely fit the critical eigenfunctions.

To see why it's unsurprising that the adaptive time-stepper should stabilize at the stability threshold, we first look at the local truncation error for the SBDF2 scheme, applied to the ODE $u_t = f(u) + g(u)$, close to a steady state u_{ss}:

$$LTE = u^{n+1} - u(t_{n+1}) = d^{n+1} - d(t_{n+1})$$
$$= \left(\frac{2}{3} d'''(t_n) - g''(u(t_n)) d'(t_n)^2 - g'(u(t_n)) d''(t_n) \right) dt^3 + O(dt^4)$$

where $u(t_n) = u^n$, $d(t_{n+1}) = u(t_{n+1}) - u_{ss}$, and $d^{n+1} = u^{n+1} - u_{ss}$. If the steady state is asymptotically stable then we assume the deviation is decaying exponentially
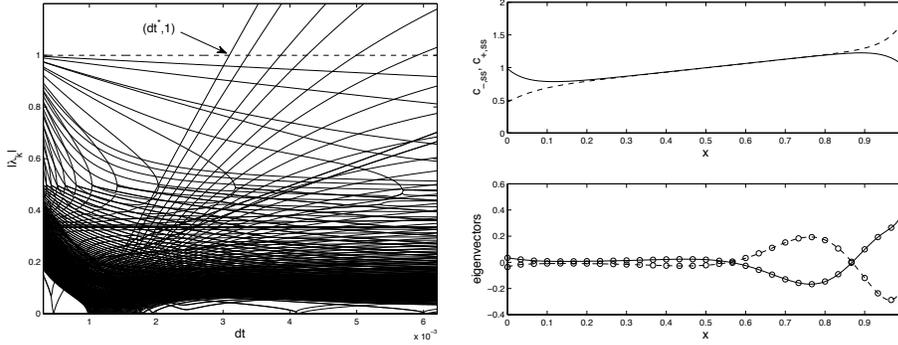
Fig. 3: PNP-FBV system (10)–(17) with constant voltage $v(t) = 2$. $\epsilon = .05$ and all other physical parameters equal 1. The mesh is uniform: $dx = 1/90$. Left plot: The magnitudes of the 364 eigenvalues are plotted versus dt. The largest magnitude branch crosses at $dt^* = .003094$. Top plot on right: The steady state profiles — the solid line is $\mathbf{c}_{p,ss}$ and the dashed line is $\mathbf{c}_{m,ss}$. Bottom plot on right: One real-valued eigenvalue goes unstable ($\lambda = -1$). The corresponding eigenvectors are plotted — the solid line is the unstable eigenvector for $\mathbf{c}_p$ and the dashed line is the unstable eigenvector for $\mathbf{c}_m$. Both eigenvectors have been chosen to have $l-2$ norm 1. The open circles denote late-time deviations from the steady states, as computed using the VSSBDF2 adaptive time-stepper. The deviations have been normalized to have l^2 norm 1; only a third of the $N = 91$ data points are plotted for tidiness.

in time: $d(t) = C \exp(-\lambda t)$. The local truncation error can be bounded

$$|C| \, \alpha \, e^{-\lambda t_n} dt^3 + O(dt^4) \leq |LTE| \leq |C| \, \beta \, e^{-\lambda t_n} dt^3 + O(dt^4)$$

where α and β are determined by λ and uniform bounds on g' and g'' near u_{ss}. For this reason, if dt is held fixed the LTE will decay to zero. If the LTE is required to be greater than $tol - range$ then dt must grow exponentially to satisfy this. The above argument is predicated on the solution's decaying according to the underlying ODE; specifically $d(t_n) = C \exp(-\lambda \, t_n) \to 0$.

For a system of ODEs, one would consider the spectral radius of the linearized scheme. If dt is such that $|\lambda(dt)|_{\max} < 1$ then the LTE for the SBDF2 scheme will go to zero exponentially fast as $n \to \infty$. Similarly, if $|\lambda(dt)|_{\max} > 1$ then the LTE will diverge. Our simulations with the VSSBDF2 adaptive time-stepper are taking essentially-constant time-steps, dt_∞, and the LTE is staying in $[tol-range, tol+range]$ where $tol - range > 0$. This behaviour can only happen if $|\lambda(dt_\infty)|_{\max} = 1$; i.e., $dt_\infty = dt^*$.

4.2. Dependence of the stability domain on ϵ. We now consider the stability properties of the PNP-FBV system (10)–(17) for a range of values of ϵ, holding the voltage fixed. We find that for $\epsilon \in (0.107764, 0.134504)$ the instability takes the form of a pair of complex eigenvalue crossing the unit circle; for all other values we considered it was a single eigenvalue crossing at -1. Figure 4 is the analogue of Figure 3 but for a value of ϵ that results in two complex eigenvalues crossing the unit circle.

For each ϵ, we find the stability threshold dt^*. The left plot of Figure 5 presents dt^* as a function of ϵ. The open circles indicate the interval of ϵ values for which a pair of complex conjugate eigenvalues cross the unit circle. The stability threshold dt^* is a continuous function of ϵ except for a jump at $\epsilon \approx 0.134504$. Also, dt^* appears
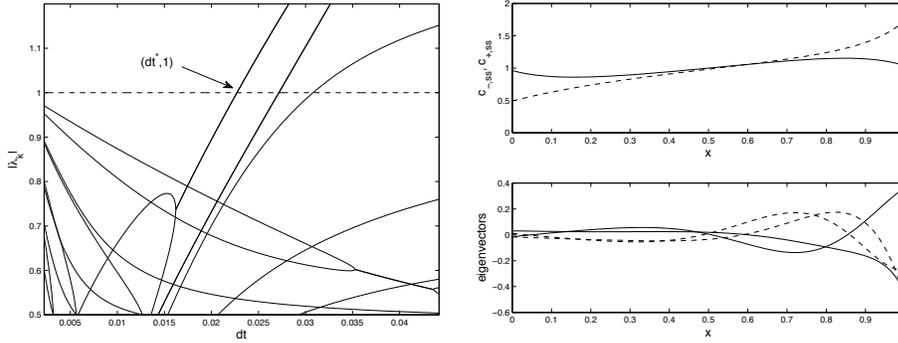
Fig. 4: PNP-FBV system (10)–(17) with constant voltage $v(t) = 2$, $\epsilon = .12$ and all other physical parameters equal 1. The mesh is uniform: $dx = 1/90$. Left plot: the magnitudes of the 364 eigenvalues are plotted versus dt; the vertical range has been truncated for a tidier plot. The largest magnitude branch crosses at $dt^* = .02271$. Top plot on right: The steady state profiles — solid line is $\mathbf{c}_{+,ss}$ and dashed line is $\mathbf{c}_{-,ss}$. Bottom plot on right: A pair of complex-valued eigenvalues go unstable $(-0.9797 \pm 0.2008\, i)$. The corresponding eigenvectors are plotted — the solid lines are the unstable eigenvectors for $\mathbf{c}_p$ and the dashed line are the unstable eigenvectors for $\mathbf{c}_m$.

to be a smooth function of ϵ except at $\epsilon \approx 0.134504$ where there's a jump in dt^* and at $\epsilon \approx 0.107764$ where there's a jump in the first derivative of dt^*.

The jump in dt^* is striking — if one were using an SBDF2 time-stepper with $dt = .025$ then this would yield a stable simulation for ϵ is close to, but slightly smaller than, the critical value of $\epsilon \approx 0.134504$ but the simulation would be unstable simulation for ϵ that is close to, but slightly larger than, this critical value. The stability of the SBDF2 simulation is not a continuous function of the parameter ϵ.

The right plot of Figure 5 compares dt_∞ as found from the VSSBDF2 adaptive time-stepper to dt^* as found from the linear stability study of the steady state. The solid line plots dt^* versus ϵ; the circles plot dt_∞. The circles align closely with the solid lines, providing compelling evidence that it is the numerical instability of the scheme near the steady state which is causing the VSSBDF2 adaptive time-stepper to stabilize its time-steps.

Figure 6 addresses the cause of the corner in the graph of $dt^*(\epsilon)$ at $\epsilon \approx .107764$. The top left figure presents a closer view of dt^* versus ϵ; there's clearly a corner in the graph. The bottom left figure presents the magnitude of the imaginary part of the eigenvalue(s) that is crossing the unit circle. For ϵ close to, but smaller than, $\epsilon = .107764$ the instability arises when a single real-valued eigenvalue crosses the unit circle through the point -1. For ϵ close to, but larger than, $\epsilon = .107764$ the instability arises when a complex conjugate pair of eigenvalues with nonzero imaginary part cross the unit circle. The figures to the right present the magnitudes of the eigenvalues as a function of dt for two values of ϵ close to $\epsilon = .107764$. The top figure is for an ϵ that is close to, but smaller than, $\epsilon = .107764$ and the top figure is for an ϵ value that is slightly larger than this critical value of ϵ. In both figures, there's a branch which denoted with a dot-dash line. This branch corresponds to a pair of complex conjugate eigenvalues; following this branch leftward and downward in the figure, one sees that it arose from the collision of two real-valued eigenvalues (there's a triple junction). In the top ($\epsilon = .107$) figure, the branch is to the right of the branch
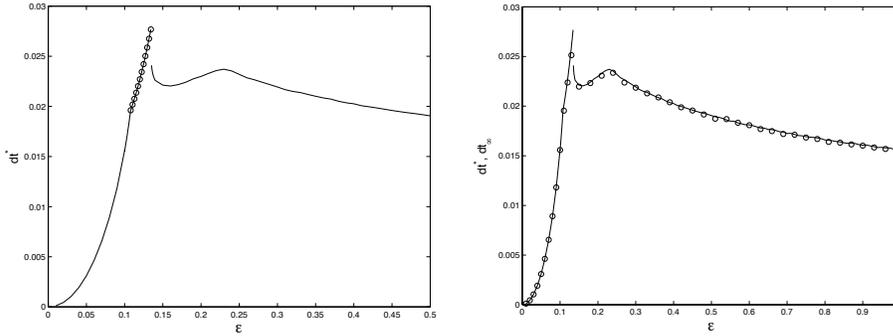
Fig. 5: PNP-FBV system (10)–(17) with parameters as for Figures 3 and 4 except for ϵ which varies. Left figure: Solid line: dt^* versus ϵ where dt^* is found from the linear stability analysis. Open circles denote ϵ values for which the instability involved a pair of complex eigenvalues crossing the unit circle: $\epsilon \in (0.107764, 0.134504)$. Right figure: Solid line: dt^* versus ϵ where dt^* is found from the linear stability analysis. Open circles: dt_∞ as found by the time-step size stabilizing in the VSSBDF2 adaptive time-stepper.

with one real eigenvalue: the complex pair of eigenvalues are not the cause of the stability threshold that corresponded to the first-crossing of the unit circle. In the bottom ($\epsilon = .109$) figure, the two branches have exchanged positions. If one views a sequence of these figures as ϵ increases from .107 to .109, one sees that both branches are moving rightward but that the branch that carries the single real eigenvalue is moving rightward at a slightly faster speed; as a result it overtakes the branch that carries the complex pair of eigenvalues.
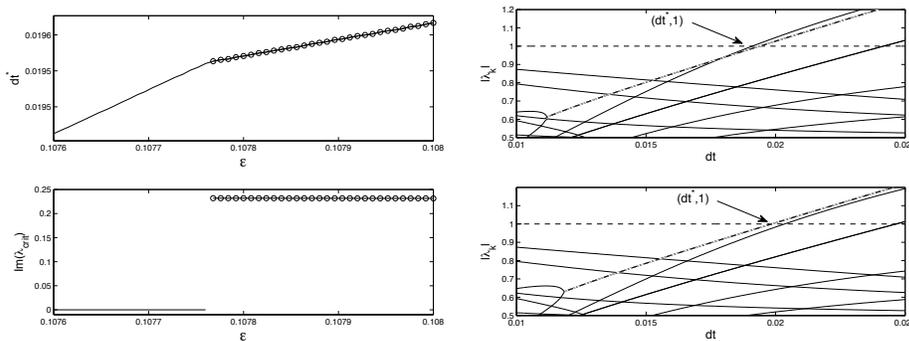
Fig. 6: PNP-FBV system (10)–(17) with parameters as for Figures 3 and 4. Left figure: Top plot presents dt^* versus ϵ where dt^* is found from the linear stability analysis. There's a corner in the graph at approximately $\epsilon = .107764$. Bottom plot presents the magnitude of the imaginary part of the eigenvalue(s) that is crossing the unit circle. There's a jump at approximately $\epsilon = 0.107764$. Right figure: Top plot presents the magnitude of the eigenvalues of the linearized problem versus dt versus for $\epsilon = .107$. Top plot presents the magnitude of the eigenvalues of the linearized problem versus dt versus for $\epsilon = .109$.

Figure 7 is the analogue of Figure 6; it addresses the cause of the jump in the graph of $dt^*(\epsilon)$ at $\epsilon \approx 0.134504$. From the figure in the left, we see that there's a jump in the stability threshold dt^* and that the eigenvalues switch from a complex

conjugate pair to a single real eigenvalue as ϵ increases through the critical value. The upper right plot presents the magnitude of the eigenvalues for a value of ϵ that is slightly smaller than the critical value and the bottom right plot presents them for a value that is slightly larger. In both plots, we see that the triple point, where the branch carrying the complex pair of eigenvalues emerges from the intersection of two branches carrying single real eigenvalues, is close to the dashed line at height 1. In the upper plot, we see that the upper branch (before the triple point) is below the dashed line — the first eigenvalues to cross the unit circle are the complex pair, for a larger value of dt^*. However, as ϵ increases, this upper branch (before the triple point) moves upwards and it reaches the dashed line when $\epsilon \approx 0.134504$; at this value of ϵ the stability threshold dt^* jumps downwards. After this critical value of ϵ, the stability threshold is due to a single real eigenvalue crossing the unit circle at -1.
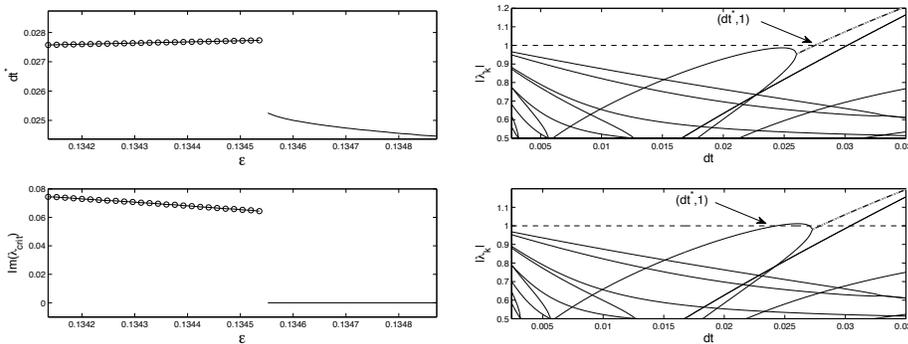
Fig. 7: PNP-FBV system (10)–(17) with parameters as for Figures 3 and 4. Left figure: Top plot presents dt^* versus ϵ where dt^* is found from the linear stability analysis. There's a jump at approximately $\epsilon \approx .134504$. Bottom plot presents the magnitude of the imaginary part of the eigenvalue(s) that is crossing the unit circle. Right figure: Top plot presents the magnitude of the eigenvalues of the linearized problem versus dt versus for $\epsilon = .134$. Top plot presents the magnitude of the eigenvalues of the linearized problem versus dt versus for $\epsilon = .135$.

4.3. Dependence of stability domain on spatial discretization. We next consider the effect of the mesh on the stability threshold. Before considering the PNP-FBV model, we first consider two simple examples. Consider the diffusion equation

$$u_t = D_1 \, u_{xx} + D_2 \, u_{xx}, \qquad \text{with} \qquad D_1 + D_2 > 0$$

subject to homogeneous Dirichlet boundary conditions: $u(0,t) = u(1,t) = 0$. This can be time-stepped using SBDF2 (1) with $D_1 u_{xx}$ handled implicitly and $D_2 u_{xx}$ handled explicitly. Using the linear stability analysis presented in Appendix A, if $D_1 < 3 \, D_2$ then there is a stability threshold on the time-step, the scheme is stable if $dt < dt^* = 4/(|\lambda_N|(3D_2 - D_1))$ and unstable otherwise. Here, λ_N is the most-negative eigenvalue of the matrix that approximates $\frac{\partial^2}{\partial x^2}$ with homogeneous Dirichlet boundary conditions: D_{xx}. For a uniform mesh, $\lambda_N \to -\infty$ like N^2 and so the stability constraint, dt^*, decays to zero like $1/N^2$. If the time-step violates the stability constraint, the fastest growing eigenvector is the one that approximates the highest-frequency resolvable eigenfunction: $\sin((N-2)\pi x)$. On the other hand, consider the

diffusion equation with a sink

$$u_t = D_1 u_{xx} - \frac{1}{\epsilon^2} u, \qquad \text{with} \qquad D_1 > 0$$

subject to homogeneous Dirichlet boundary conditions: $u(0,t) = u(1,t) = 0$. This can be time-stepped using SBDF2 (1) with $D_1 u_{xx}$ handled implicitly and u/ϵ^2 handled explicitly. The linear stability analysis presented in Appendix A implies that if $\epsilon^2 D_1 |\lambda_1| < 3$ then there is a stability threshold on the time-step; the scheme is stable if $dt < dt^* = 4/(D_1\lambda_1 + 3/\epsilon^2)$ and unstable otherwise. Here, λ_1 is the largest nonzero eigenvalue of D_{xx}. For a uniform mesh, $\lambda_1 \to -\pi^2$ as $N \to \infty$. From this, we see that the stability constraint, dt^*, does depend on the mesh but it does so in a gentle manner: it converges to a positive number as the mesh is refined. If the time-step violates the stability constraint, the fastest growing mode is the eigenvector that approximates the low frequency eigenfunction $\sin(\pi x)$.

Returning to the PNP-FBV system, the plot in the left of Figure 8 presents stability thresholds for four different meshes: one uniform mesh and three piecewise uniform meshes that have a finer mesh near $x = 0, 1$. The stability threshold depends on the mesh in a mild manner. This is not surprising given that the critical eigenmodes of the linearized scheme presented in the bottom right plots of Figures 3 and 4 do not appear to have structures that need significant spatial resolution.
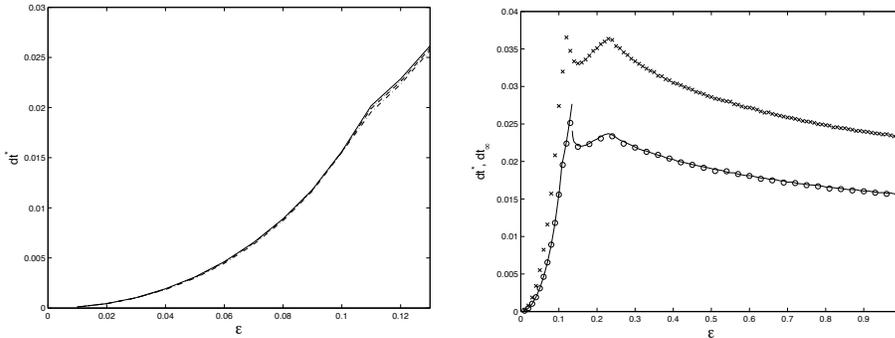
Fig. 8: PNP-FBV system (10)–(17) with constant voltage ($v(t) = 2$) ϵ varies, all other physical parameters set to 1. Left plot: For four different meshes, dt^* is computed and plotted against ϵ. Solid line: uniform mesh with $dx = 1/90$. Dot-dashed line: piecewise uniform mesh with $dx = 1/150$ in $[0, 1/10]$ and $[9/10, 1]$ and $dx = 4/75$ elsewhere. Dashed line: piecewise uniform mesh with $dx = 1/300$ in $[0, 1/10]$ and $[9/10, 1]$ and $dx = 2/75$ elsewhere. Dotted line: piecewise uniform mesh with $dx = 1/450$ in $[0, 1/10]$ and $[9/10, 1]$ and $dx = 4/225$ elsewhere. Right plot: The parameters are as in Figure 5. Solid line: dt^* versus ϵ where dt^* is found from the linear stability analysis. Open circles: dt_∞ as found by the time-step size stabilizing in the VSSBDF2 adaptive time-stepper with no Richardson extrapolation step. X-marks: dt_∞ as found by the VSSBDF2 adaptive time-stepper with a Richardson extrapolation step.

4.4. Effect of Richardson Extrapolation. Richardson extrapolation is commonly used to increase the accuracy of time-stepping, however we have found much discussion of its possible effect on the stability of the scheme. In [25], the authors present an analysis of forward Euler time-stepping for a problem of interest. They demonstrate analytically how Richardson extrapolation affects the stability threshold between conditionally stable and unconditionally stable. Below, we computationally

study the effect of Richardson extrapolation on the conditional stability for the PNP-FBV system. We also present an analytical and computational study of the effect of Richardson extrapolation on SBDF2 simulations of the heat equation.

As described in Section 2, the adaptive time-stepper computes two approximations of the solution at time t^{n+1} — the "coarse" approximation $(\mathbf{u}_c^{n+1})$ and the "fine" approximation $(\mathbf{u}_f^{n+1})$ — and uses these to approximate the local truncation error. The solution at time t^{n+1} is then either taken to equal $\mathbf{u}_c^{n+1}$ (see (4)) or Richardson extrapolation is used to blend $\mathbf{u}_c^{n+1}$ and $\mathbf{u}_f^{n+1}$ (see (5)). In the simulations presented up to this point, we did not use Richardson extrapolation in defining $\mathbf{u}^{n+1}$. That said, we find that when using the VSSBDF2 adpative time-stepper with Richardson extrapolation (5) on the PNP-FBV system, the observed behaviour is like that when Richardson extrapolation was not used: the time-steps stabilized at a value dt_∞.

Both $\mathbf{u}_c^{n+1}$ and $\mathbf{u}_f^{n+1}$ both satisfy time-stepping equations like (2) and (19) but $\mathbf{u}^{n+1}$ doesn't. Our linear stability analysis is built upon the solutions satisfying (19) and so its predictions only apply to the VSSBDF2 adaptive time-stepper when Richardson extrapolation is not used. We haven't found an analogue of (2) that $u^{n+1} = \alpha u_c^{n+1} + \beta u_f^{n+1}$ satisfies. Given such an analogue, we could perform the linear stability analysis to determine a stability threshold dt^* to compare to the dt_∞ found by the VSSBDF2 adaptive time-stepper when Richardson extrapolation is used. Because dt_∞ and dt^* agree closely when Richardson extrapolation isn't used; we use dt_∞ as a proxy for the stability threshold dt^* when Richardson extrapolation is used.

The plot in the right of Figure 8 presents dt^* and dt_∞ where dt_∞ is found using two different implementations of the VSSBDF2 adaptive time-stepper. The open circles denote dt_∞ as found by the VSSBDF2 adaptive time-stepper with no Richardson extrapolation (4). The crosses denote dt_∞ as found by the VSSBDF2 adaptive time-stepper with Richardson extrapolation (5). We see that when Richardson extrapolation is used dt_∞ is larger than dt^*, sometimes markedly so, and for this reason the simulations finish more quickly when Richardson extrapolation is used. It's also striking that the plot of dt_∞ versus ϵ when Richardson extrapolation is used has a very similar shape to the plot of the data when it isn't used.

For the PNP-FBV system, we found that using Richardson extrapolation as part of the adaptive time-stepper leads to greater stability. However, this is problem dependent. We found that if one repeats this experiment for the diffusion equation $u_t = D_1 u_{xx} + D_2 u_{xx}$, with solutions simulated as described in §§4.3, then for some choices of D_1 and D_2 using Richardson extrapolation in the VSSBDF2 adaptive time-stepper leads to less stability: for some parameter choices dt_∞ is smaller when Richardson extrapolation is used. There are even parameter choices for which the SBDF2 time-stepper is unconditionally stable (and so the VSSBDF2 adaptive time-stepper would with no Richardson extrapolation will have dt increase to dt_{max}) but if one uses Richardson extrapolation in the adaptive time-stepper then a limiting step size dt_∞ is observed. This suggests that for such parameter choices, Richardson extrapolation has changed the underlying time-stepping scheme from unconditionally stable to conditionally stable.

We performed a range of simulations in which D_1 and N were held fixed and D_2 was varied while keeping $3 D_2 > D_1$ so that the SBDF2 scheme is conditionally stable. Each simulation was run sufficiently long that dt_∞ could be identified. In Figure 9, $1/dt_\infty$ is plotted against D_2; the open circles indicate $1/dt_\infty$ when no Richardson extrapolation (4) is used in the adaptive time-stepper and the X-marks indicate $1/dt_\infty$ when Richardson extrapolation is used (5). The solid line is the
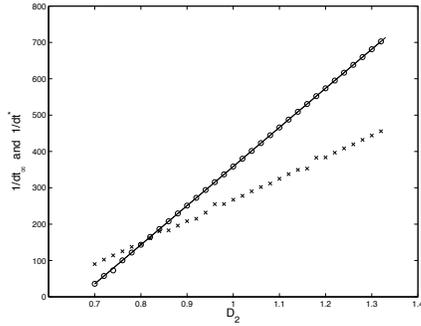
Fig. 9: $u_t = D_1 u_{xx} + D_2 u_{xx}$ with initial data $\sin(\pi x)$ and boundary conditions $u(0,t) = u(1,t) = 0$ is simulated on $[0,1]$ using VSSBDF2 adaptive time-stepper. $N = 20$, $tol = 10^{-6}$ and $range = tol/3$. D_1 is held fixed $D_1 = 2$ and D_2 takes values between $1.05 D_1/3$ and $2D_1/3$. Circles: $1/dt_\infty$ where dt_∞ is the limiting time-step size found by the VSSBDF2 adaptive time-stepper with no Richardson Extrapolation (4). Line: $1/dt^*$ versus D_2 where $dt^* = 4/(|\lambda_N|(3D_2 - D_1))$. There is no stability constraint if $D_2 \leq D_1/3$. X-marks: $1/dt_\infty$ where dt_∞ is the limiting time-step size found by the VSSBDF2 adaptive time-stepper with Richardson Extrapolation (5). Linear fit to X-marks: $|\lambda_N|(.4124 D_2 - .1130 D_1)$; there is no stability constraint if $D_2 < .2739 D_1$.

plot of $1/dt^* = |\lambda_N|(3D_2 - D_1)/4$. Fitting the X-marks using least squares, we find that $1/dt_\infty \approx |\lambda_N|(.41 D_2 - .11 D_1)$. Without Richardson Extrapolation, the scheme is has no stability constraints if $D_2 \leq D_1/3$. With Richardson Extrapolation, this unconditional stability is for a smaller range: if $D_2 \lesssim .27 D_1$. For D_2 in the (approximate) interval $(.27 D_1, .41 D_1)$ the Richardson Extrapolation is destabilizing: either it introduces a stability constraint or both schemes have a stability constraint and dt_∞ is smaller when Richardson Extrapolation is used. For $D_2 \gtrsim .41 D_1$, the Richardson Extrapolation has a stabilizing effect: dt_∞ is larger when Richardson Extrapolation is used.

4.5. Dependence of stability domain on applied voltage and on the singular perturbation parameter. We next studied how the stability threshold, dt^*, depends on the applied voltage. For this, we imposed constant voltages, with values ranging between 0 and 3. The left plot of Figure 10 presents the stability threshold, dt^*, versus ϵ for the four voltages. We see that for smaller values of ϵ, dt^* does not appear to be affected as much by the applied voltage compared to larger values of ϵ. The vertical dashed line in the upper left plot indicates $\epsilon = .5$; this was the ϵ used in the time-dependent voltage simulation presented in Figure 4 of the companion article [38]. The four intersection points of the dashed line with the graphs of dt^* are the values marked with dotted lines in Figure 4 of the companion article [38]. The lower left plot of Figure 10 suggests that dt^* may be proportional to a power of ϵ for small values of ϵ. The right plot of Figure 10 presents $\ln(dt^*)$ versus $\ln(\epsilon)$. The four graphs appear to be roughly linear but don't appear to have the same slopes. All four plots correspond to dt^* decreasing to zero slightly faster than ϵ^2, consistent with Table 2 in the companion article [38].

5. Conclusions and Future Work. In this work, we considered the Poisson-Nernst-Planck equations with generalized Frumkin-Butler-Volmer reaction kinetics at the electrodes. When the VSSBDF2 adaptive time-stepper is being used to study scenarios in which the voltage or the current are (nearly) constant for long periods
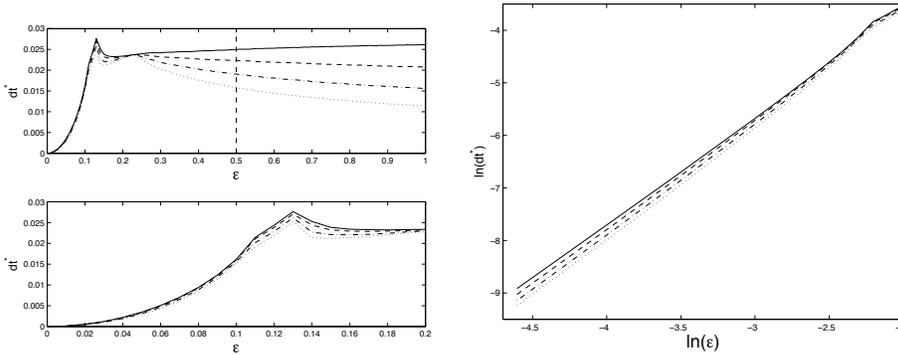
Fig. 10: PNP-FBV system (10)–(17) with parameters as for Figures 3 and 4. Four values of (constant) voltage are considered: $v(t) = 0, 1, 2$, and 3. Left figure: The threshold time-step dt^* versus ϵ for $v(t) = 0$ (solid), $v(t) = 1$ (dashed), $v(t) = 2$ (dot-dashed), and $v(t) = 3$ (dotted). Right figure: Here, $\log(dt^*)$ is plotted versus $\log(\epsilon)$ for the same voltage values and line notations. The decay is roughly like ϵ^2; fitting the data yields exponents 2.0677, 2.1145, 2.1540, and 2.1813 for constant voltages $v(t) = 0$, 1, 2, and 3 respectively.

of time, the time-step sizes stabilize to a limiting value and the computed solutions nearly converge to a steady state, but fail to do so. This behaviour is understood by linearizing the numerical scheme about the steady state. The linearized scheme is found to be conditionally stable, with a stability threshold that agrees with the time-step at which that the adaptive time-stepper stabilized. The stability domain is studied numerically and is found to have a corner and a jump discontinuity. The eigenfunctions corresponding to the critical eigenvalues are studied; the conditional stability is not related to a high-frequency instability. Using a Richardson extrapolation step in the adaptive time-stepper appears to stabilize the problem somewhat in that the limiting time-step is larger. However other systems are presented in which Richardson extrapolation can destabilize the scheme.

It would be interesting to see if one can modify the conditional stability by using information about the structure of the steady state. For example, [20] created an unconditionally stable scheme for a nonlinear diffusion equation by using bounds on the solution, although the instability being controlled was due to high frequencies.

Our methods are not restricted to the PNP-FBV system or to the VSSBDF2 adaptive time-stepper. If one is using linear multi-step method to study a system that has asymptotically stable steady states, our approach is relevant. We expect that it would generalize in a natural manner to Runge-Kutta methods as well. A natural next step would be to study the stability properties of IMEX schemes beyond steady states by considering problems that have orbitally stable special solutions, such as travelling waves, or by considering problems that have asymptotically stable special solutions, such as self-similar solutions.

Appendix A. Stability for an SBDF2 scheme.
In this section, we consider time-stepping the ODE

$$u_t = (\lambda + \alpha)\, u = \lambda\, u + \alpha\, u \quad \text{with } \alpha, \lambda \neq 0,$$

using the SBDF2 scheme (1) with $f(u) = \alpha\, u$ and $g(u) = \lambda\, u$.

Solutions of the ODE decay to the steady-state solution $u_\infty(t) = 0$ as $t \to \infty$. The

numerical stability question is: under what conditions do solutions of the numerical scheme decay to 0 as $n \to \infty$? The general solution of (1) is $u^n = c_1 \rho_+^n + c_2 \rho_-^n$ where

$$(25) \qquad \rho_\pm(dt) = \frac{-2 - 2\alpha\, dt \pm \sqrt{1 + 2\, dt(1 + 2\,\alpha\, dt)\,(\alpha + \lambda)}}{-3 + 2\, dt\, \lambda}.$$

One can verify that $\rho_-(dt)$ is the essential root and $\rho_+(dt)$ is the spurious root: $\rho_+(dt) \to 1/3$ and $\rho_-(dt) \to 1$ as $dt \to 0$. Also, $\rho'_-(0) = \lambda + \alpha < 0$ and so $\rho_-(dt) < 1$ for dt close to 0. By continuity of $\rho_\pm$, it follows that $-1 < \rho_\pm(dt) < 1$ for small dt.

If dt is such that $|\rho_+(dt)| > 1$ or $|\rho_-(dt)| > 1$, then the scheme will be numerically unstable for that time-step size.

First, we assume that $\rho_\pm(dt)$ are real-valued and seek solutions of $\rho_\pm(dt) = 1$. Doing this, we find that $\rho_\pm(dt) = 1$ is impossible for $dt > 0$. If $3\alpha < \lambda$ then $\rho_+(dt^*) = -1$, where $dt^* = 4/(\lambda - 3\alpha)$, and $\rho_-(dt) = -1$ has no solutions on $(0, \infty)$. If $3\alpha \geq \lambda$ then $\rho_\pm(dt) = -1$ has no solutions on $(0, \infty)$.

If, on the other hand, $\rho_\pm(dt)$ are complex valued then they are complex conjugates. A short calculation for complex-valued $\rho_\pm(dt)$ yields

$$|\rho_\pm(dt)|^2 = \frac{1 + 2\alpha\, dt}{3 - 2\lambda\, dt} = 1 \quad \Longleftrightarrow \quad dt = \frac{1}{\lambda + \alpha}.$$

By assumption, $\lambda + \alpha < 0$ and so there is no $dt > 0$ such that $\rho_\pm(dt)$ is complex-valued and $|\rho_\pm(dt)| = 1$.

The parameters α and λ determine whether the roots $\rho_\pm$ are real or complex valued; the roots of the discriminant in (25) are

$$(26) \qquad dt_\pm = -\frac{1}{4\alpha}\left(1 \pm \sqrt{1 - \frac{4\alpha}{\lambda + \alpha}}\right).$$

We consider the following four cases for the numerical stability study:

Case 1: $\alpha > 0$ The roots $\rho_\pm$ are real-valued on $[0, dt_-]$ and are complex-valued on (dt_-, ∞). Because $\lambda + \alpha < 0$, it's impossible to have both $\alpha > 0$ and $3\alpha < \lambda$. It follows that $\rho_\pm(dt) = -1$ has no solutions on $(0, dt_-]$. We already know that $\rho_\pm(dt) = 1$ has no solutions on $(0, dt_-]$ and that there are no solutions of $|\rho_\pm(dt)| = 1$ on (dt_-, ∞). The roots $\rho_\pm(dt)$ are continuous on $[0, \infty)$; it follows that $|\rho_\pm(dt)| < 1$ on $[0, \infty)$. Indeed, $-1 < \rho_\pm(dt) < 1$ for small values of dt and it would be impossible for $|\rho_-(dt)| > 1$ or $|\rho_+(dt)| > 1$ for any subsequent value because, by continuity, this would require $|\rho_-|$ or $|\rho_+|$ to equal 1 at some prior value.

Case 2: $\alpha < 0$ and $3\alpha > \lambda$ The roots $\rho_\pm$ are real-valued on $[0, dt_-]$, complex-valued on (dt_-, dt_+), and real-valued on $[dt_+, \infty)$. Because $\lambda < 3\alpha$, $\rho_\pm(dt) = -1$ has no solutions on $(0, dt_-] \cup [dt_+, \infty)$. We already know that $\rho_\pm(dt) = 1$ has no solutions on $(0, dt_-] \cup [dt_+, \infty)$ and that there are no solutions of $|\rho_\pm(dt)| = 1$ on (dt_-, dt_+). By the same continuity arguments as in Case 1, it follows that $|\rho_\pm(dt)| < 1$ on $[0, \infty)$.

Case 3: $\alpha < 0$ and $3\alpha = \lambda$ The roots $\rho_\pm(k)$ simplify to $1/3$ and $(-3 - 6\alpha k)/(-3 + 6\alpha k)$. They are continuous on $\mathbb{R}$ and take values in $(-1, 1)$.

Case 4: $\alpha < 0$, $3\alpha < \lambda$, and $\lambda < 0$ The roots $\rho_\pm$ are real-valued and continuous on $[0, \infty)$. There are no solutions of $\rho_-(dt) = -1$ or $\rho_\pm(dt) = 1$. There is exactly one

solution of $\rho_+(dt) = -1$; at $dt^* = 4/(\lambda - 3\alpha)$. By the same continuity arguments as in Case 1, it follows that $-1 \leq \rho_\pm(dt) < 1$ on $(0, dt^*]$ and $\rho_+(dt) < -1$ on (dt^*, ∞).

Case 5: $\alpha < 0$, $3\alpha < \lambda$, and $\lambda > 0$ The roots $\rho_\pm$ are real-valued and continuous on $[0, 3/(2\lambda))$. As dt increases to $3/(2\lambda)$, $\rho_+(dt)$ diverges to $-\infty$. By the same arguments as in Case 3, $-1 \leq \rho_\pm(dt) < 1$ on $(0, dt^*]$ and $\rho_+(dt) < -1$ on $(dt^*, 3/(2\lambda))$.

In summation, a numerical instability is only possible in Cases 4 and 5: if $\alpha < 0$ and $3\alpha < \lambda$. In this case, $dt^* = 4/(\lambda - 3\alpha)$ is the time-step limit.

REFERENCES

[1] P. M. Biesheuvel and M. Z. Bazant. Nonlinear dynamics of capacitive charging and desalination by porous electrodes. *Phys. Rev. E*, 81(3):031502, 2010.

[2] P. M. Biesheuvel, Yeqing Fu, and M. Z. Bazant. Diffuse charge and Faradaic reactions in porous electrodes. *Phys. Rev. E*, 83(6):061507, 2011.

[3] P. M. Biesheuvel, Yeqing Fu, and M. Z. Bazant. Electrochemistry and capacitive charging of porous electrodes in asymmetric multicomponent electrolytes. *Russian J. Electrochem.*, 48(6):580–591, 2012.

[4] P. B. Peters, R. van Roij, Martin Z. Bazant, and P. M. Biesheuvel. Analysis of electrolyte transport through charged nanopores. *Phys. Rev. E*, 93:053108, 2016.

[5] Ian Streeter and Richard G. Compton. Numerical simulation of potential step chronoamperometry at low concentrations of supporting electrolyte. *J. Phys. Chem. C*, 112(35):13716–13728, 2008.

[6] Richard G. Compton and Craig E. Banks. *Understanding Voltammetry.* Imperial College Press, London, 2011.

[7] E. Victoria Dydek and Martin Z. Bazant. Nonlinear dynamics of ion concentration polarization in porous media: The leaky membrane model. *AlChE Journal*, 59(9):3539–3555, 2013.

[8] Victor V. Nikonenko, Natalia D. Pismenskaya, Elena V. Belova, Philippe Sistat, Patrice Huguet, Gérald Pourcelly, and Christian Larchet. Intensive current transfer in membrane systems: Modelling, mechanisms and application in electrodialysis. *Adv. Colloid Interface Sci.*, 160(1):101–123, 2010.

[9] A. Yaroshchuk. Over-limiting currents and deionization shocks in current-induced polarization: local equilibrium analysis. *Adv. Colloid Interface Sci.*, 183:68–81, 2012.

[10] Martin Z. Bazant and Todd M. Squires. Induced-charge electrokinetic phenomena. *Current Opinion in Colloid & Interface Science*, 15(3):203–213, 2010.

[11] Martin Z. Bazant, Mustafa Sabri Kilic, Brian D. Storey, and Armand Ajdari. Towards an understanding of induced-charge electrokinetics at large applied voltages in concentrated solutions. *Advances in colloid and interfacial science*, 152(1):48–88, 2009.

[12] Martin Z. Bazant, Brian D. Storey, and Alexei A. Kornyshev. Double layer in ionic liquids: Overscreening versus crowding. *Phys. Rev. Lett.*, 106(4):046102, 2011.

[13] Alexei A. Kornyshev. Double-layer in ionic liquids: Paradigm change? *J. Phys. Chem. B*, 111:5545–5557, 2007.

[14] Martin Z. Bazant, Kevin T. Chu, and B. J. Bayly. Current-voltage relations for electrochemical thin films. *SIAM J. Appl. Math.*, 65(5):1463–1484, 2005.

[15] Kevin T. Chu and Martin Z. Bazant. Electrochemical thin films at and above the classical limiting current. *SIAM J. Appl. Math.*, 65(5):1485–1505, 2005.

[16] P. M. Biesheuvel, M. van Soestbergen, and M. Z. Bazant. Imposed currents in galvanic cells. *Electrochemica Acta*, 54:4857–4871, 2009.

[17] P. Maarten Biesheuvel, Alejandro A. Franco, and Martin Z. Bazant. Diffuse charge effects in fuel cell membranes. *J. Electrochem. Soc.*, 156(2):B225–B233, 2009.

[18] Alpha A Lee, Svyatoslav Kondrat, Gleb Oshanin, and Alexei A Kornyshev. Charging dynamics of supercapacitors with narrow cylindrical nanopores. *Nanotechnology*, 25(31):315401, 2014.

[19] Rodolfo R. Rosales, Benjamin Seibold, David Shirokoff, and Dong Zhou. Unconditional

Stability for Multistep ImEx Schemes: Theory. *SIAM Journal on Numerical Analysis*, 55(5):2336–2360, 2017.

[20] Benjamin Seibold, David Shirokoff, and Dong Zhou. Unconditional stability for multistep ImEx schemes: Practice. *Journal of Computational Physics*, 376:295–321, JAN 1 2019.

[21] Georgios Akrivis and Fotini Karakatsani. Modified implicit-explicit BDF methods for nonlinear parabolic equations. *BIT*, 43(3):467–483, 2003.

[22] Georgios Akrivis, Ohannes Karakashian, and Fotini Karakatsani. Linearly implicit methods for nonlinear evolution equations. *Numer. Math.*, 94(3):403–418, 2003.

[23] Jim Douglas, Jr. and Todd Dupont. Alternating-direction Galerkin methods on rectangles. In *Numerical Solution of Partial Differential Equations, II (SYNSPADE 1970) (Proc. Sympos., Univ. of Maryland, College Park, Md., 1970)*, pages 133–214. Academic Press, New York, 1971.

[24] Thomas Y. Hou, John S. Lowengrub, and Michael J. Shelley. Removing the stiffness from interfacial flows with surface tension. *J. Comput. Phys.*, 114(2):312–338, 1994.

[25] Laurent Duchemin and Jens Eggers. The Explicit-Implicit-Null method: Removing the numerical instability of PDEs. *Journal of Computational Physics*, 263:37–52, APR 15 2014.

[26] Assyr Abdulle and Alexei A. Medovikov. Second order Chebyshev methods based on orthogonal polynomials. *Numer. Math.*, 90(1):1–18, 2001.

[27] Jaemin Shin, Hyun Geun Lee, and June-Yub Lee. Unconditionally stable methods for gradient flow using convex splitting Runge-Kutta scheme. *J. Comput. Phys.*, 347:367–381, 2017.

[28] G. Izzo and Z. Jackiewicz. Highly stable implicitexplicit rungekutta methods. *Applied Numerical Mathematics*, 113:71–92, 2017.

[29] Aly-Khan Kassam and Lloyd N. Trefethen. Fourth-order time-stepping for stiff PDEs. *SIAM J. Sci. Comput.*, 26(4):1214–1233, 2005.

[30] Lili Ju, Jian Zhang, Liyong Zhu, and Qiang Du. Fast explicit integration factor methods for semilinear parabolic equations. *J. Sci. Comput.*, 62(2):431–455, 2015.

[31] Paul A. Milewski and Esteban G. Tabak. A pseudospectral procedure for the solution of nonlinear wave equations with examples from free-surface flows. *SIAM J. Sci. Comput.*, 21(3):1102–1114, 1999.

[32] Michael L. Minion. Semi-implicit spectral deferred correction methods for ordinary differential equations. *Commun. Math. Sci.*, 1(3):471–500, 2003.

[33] Oscar P. Bruno and Edwin Jimenez. Higher-Order Linear-Time Unconditionally Stable Alternating Direction Implicit Methods for Nonlinear Convection-Diffusion Partial Differential Equation Systems. *Journal of Fluids Engineering — Transactions of the ASME*, 136(6), JUN 2014.

[34] Oscar P. Bruno and Mark Lyon. High-order unconditionally stable FC-AD solvers for general smooth domains. I. Basic elements. *J. Comput. Phys.*, 229(6):2009–2033, 2010.

[35] Mark Lyon and Oscar P. Bruno. High-order unconditionally stable FC-AD solvers for general smooth domains. II. Elliptic, parabolic and hyperbolic PDEs; theoretical considerations. *J. Comput. Phys.*, 229(9):3358–3381, 2010.

[36] David Yan, Martin Z. Bazant, P. M. Biesheuvel, Mary C. Pugh, and Francis P. Dawson. Theory of linear sweep voltammetry with diffuse charge: Unsupported electrolytes, thin films, and leaky membranes. *Phys. Rev. E*, 95:033303, Mar 2017.

[37] David Yan. *Macroscopic Modeling of a One-Dimensional Electrochemical Cell Using the Poisson-Nernst-Planck Equations*. PhD thesis, University of Toronto, 2017.

[38] David Yan, Mary C. Pugh, and Francis P. Dawson. Adaptive time-stepping schemes for the solution of the Poisson-Nernst-Planck equations, 2019.

[39] D. Wang and S. J. Ruuth. Variable step-size implicit-explicit linear multistep methods for time-dependent partial differential equations. *Journal of Computational Mathematics*, 26(6):838–855, 2008.

[40] J. Rosam, P. K. Jimack, and A. Mullis. A fully implicit, fully adaptive time and space discretisation method for phase-field simulation of binary alloy solidification. *J. Comput. Phys.*, 225(2):1271–1287, 2007.

[41] U. M. Ascher, S. J. Ruuth, and B. T. R. Wetton. Implicit-explicit methods for time-dependent partial differential equations. *SIAM Journal on Numerical Analysis*, 32(3):797–823, 1995.

[42] Martin Z. Bazant. Theory of chemical kinetics and charge transfer based on nonequilibrium thermodynamics. *Accounts of Chemical Research*, 46(5):1144–1160, 2013.

[43] M. van Soestbergen. Frumkin-Butler-Volmer theory and mass transfer in electrochemical cells. *Russian Journal of Electrochemistry*, 48(6):570–579, 2012.

[44] A. A. Moya, J. Castilla, and J. Horno. Ionic transport in electrochemical cells including electrical double-layer effects. A network thermodynamics approach. *J. Phys. Chem.*, 99:1292–1298, 1995.

[45] M. van Soestbergen, P. M. Biesheuvel, and M. Z. Bazant. Diffuse-charge effects on the transient response of electrochemical cells. *Phys. Rev. E*, 81(2):1–13, 2010.

[46] RJ LeVeque. *Finite Difference Methods for Ordinary and Partial Differential Equations: Steady-State and Time-Dependent Problems*. SIAM, Philadelpha, 2007.