Redox reactions with empirical potentials: atomistic battery discharge simulations

Wolf B. Dapp[1, a)] and Martin H. Müser[1, 2, b)]

[1)]*Jülich Supercomputing Centre, Institute for Advanced Simulation, FZ Jülich, 52425 Jülich, Germany*

[2)]*Dept. of Materials Science and Engineering, Universität des Saarlandes, Saarbrücken, Germany*

(Dated: 4 June 2013)

Batteries are pivotal components in overcoming some of today's greatest technological challenges. Yet to date there is no self-consistent atomistic description of a complete battery. We take first steps toward modeling of a battery as a whole microscopically. Our focus lies on phenomena occurring at the electrode-electrolyte interface which are not easily studied with other methods. We use the redox split-charge equilibration (redoxSQE) method that assigns a discrete ionization state to each atom. Along with exchanging partial charges across bonds, atoms can swap integer charges. With redoxSQE we study the discharge behavior of a nano-battery, and demonstrate that this reproduces the generic properties of a macroscopic battery qualitatively. Examples are dependence of the battery's capacity on temperature and discharge rate, as well as performance degradation upon recharge.

I. INTRODUCTION

Batteries have been the focus of intense scrutiny in recent years. This research is driven by specialized requirements for energy storage, e.g., capabilities for high current for automotive purposes, high capacity batteries to buffer disparities in supply and demand for the power grid, and high energy density for batteries used in portable electronic devices.[1–4] Much effort is directed toward studying rechargeable lithium batteries,[5] as that is the preferred base material due to the high energy densities achievable, its good cyclability properties, a high working voltage, and its abundance in the Earth's crust.

Numerical modeling flanks experimental work in the goal to optimize batteries. Traditionally, macro-homogeneous modeling uses diffusive processes through porous material to describe batteries on a mesoscale (i.e., smaller than the electrode, larger than a molecule), and adds charge and mass balance equations as well as transfer kinetics across the boundary surfaces.[6–19] Such simulations very successfully reproduce the macroscopic behavior of batteries, and can be used to optimize parameters such as electrode thickness. Mesoscopic porous electrode models form the basis for the Li-ion battery module for instance in the commercial COMSOL multi-physics simulation package. However, some of the underlying assumptions of this approach can be considered "uncertain at best."[19]

Mesoscale modeling requires constitutive equations. It is possible to input a great many effects based on parameterized experimental data into such descriptions. Those include, but are not limited to, ionic and electronic conductivities, specific surface areas, tortuosities, porosi-ties, activity coefficients, transference numbers, concentrations, diffusion coefficients, current densities, electrochemical potential, reaction rate constants, solid electrolyte interface transport processes, contact resistances between active material and current collector, and mechanical properties such as strain tensors, elastic moduli or fracture strengths.[8,17,19–21] Unfortunately, this parameterizability which makes the approach suitable to describe real batteries also limits its general predictive power, especially on the nanoscale.

In recent molecular dynamics (MD) works, on the other hand, much effort has been devoted to dealing with certain aspects of battery design in detail. Examples are plentiful and include improving the intercalation of lithium into graphite, or calculating the transport properties of lithium through various electrode materials.[22–31] In contrast to mesoscale models, MD simulations tackle micro- and nano-physical aspects of the problem (typically only considering half-cells), but cannot make macroscopic predictions on, for instance, how the voltage during discharge changes under the influence of various control parameters such as temperature, discharge current, or changes of cell geometry.

Standard MD methods — based on either conventional *ab initio* density functional theory (DFT) or traditional *charge equilibration* (QE)[32] — cannot be used to model a battery as a whole, because those seek to equalize the chemical potential. However, the difference in the chemical potential between two electrodes is precisely what drives charge transport in a battery. Furthermore, both approaches are ill-suited to model history-dependent effects. The reason is that they carry out a unique energy minimization based on instantaneous nuclear positions. The electron transfer process during a redox reaction brings about a quasi-discontinuous change of the electronic state, modifying all molecular orbitals,[33] while the atomic configuration remains virtually unaltered.

Time-dependent DFT (TDDFT)[34–36] can elucidate

[a)]Electronic mail: w.dapp@fz-juelich.de

[b)]Electronic mail: martin.mueser@mx.uni-saarland.de

history-dependent effects, and could thus in principle model batteries, but is too expensive computationally to treat more than a few excited states. Similarly, long simulation times are prohibitively costly even with current supercomputers. Therefore, the application of TDDFT, and similarly *ab initio* MD, to electrochemical processes with regards to batteries has been limited.[37,38]

Recently, it was proposed that the limitations of MD calculations can be mitigated by introducing the oxidation state as a time-dependent variable, which needs to be subjected to dynamics.[39–41] In addition to exchanging partial charges as in the standard *split charge equilibration* (SQE) method,[42] atoms can then change their ionization state (i.e., participate in a redox reaction) by swapping integer charges across a bond (integer charge transfer – ICT). Hereafter the method is referred to as redoxSQE.

In a previous work,[41] we applied redoxSQE to case studies of contact electrification between two clusters of ideal metals and ideal dielectrics, respectively. If two initially neutral clusters with differing electron affinities are brought into contact, they will exchange charge. After separation, some portion of the transferred charge does not flow back, generating a remnant electric field that was not present before the contact. Neither conventional QE methods nor (non-time-dependent) DFT can capture this history-dependence. RedoxSQE in contrast, successfully produces charge hysteresis effects during approach and retraction, despite identical atomic positions.

In this paper we apply the redoxSQE method to a more complex problem. We want to bridge the gap between the mesoscale and highly accurate (DFT, accurate force field) approaches, and model a nano-battery. If properly parameterized, redoxSQE can be used to model the microphysics at both electrode-electrolyte interfaces, including their structural evolution and changing morphology, as well as battery performance degradation. At this stage, the simulations are meant to serve as proof-of-concept, rather than emulate any real system or produce new quantitative insights. However, even at its present qualitative level, our model reproduces generic features of battery discharge. We believe that redoxSQE can be parameterized to describe real materials quantitatively, because unmodified SQE combined with REBO force fields has yielded good agreement with experimental and DFT results (heats of formation of isolated molecules, radial distribution functions for water and ethanol, and energies of oxygenated diamond surfaces) for systems in which each element had a well-defined oxidation state.[43,44]

This paper is structured as follows. We outline the method in Section II below. We also introduce the additional parameters and procedures not covered in Ref.[41], which describes the method in greater detail. In Section IV of this paper, we present the results attained by varying both internal and external parameters, and compare the outcome to generic properties of macroscopic batteries. We close with a discussion and summary of our findings in Section V.

II. METHOD

This section briefly outlines the numerical methods (redoxSQE) used in this study, and the parameters involved. For a more detailed description we refer the reader to Ref.[41]. For comparisons with DFT-based results see the work by Verstraelen *et al.*,[40,45] whose SQE+Q^0 is similar in spirit to redoxSQE, but applies charge constraints to fragments of molecules rather than to alter the oxidation state of individual atoms.

We implement molecular dynamics with a long-range potential due to fractional charges ("split charges"), as discussed in Ref.[42] We add in the modification proposed in Ref.:[39]

$$V_{\text{total}} = V_{\text{short}} + V_{\text{long}}, \tag{1}$$

$$V_{\text{long}} = V_{\text{C}} + \sum_i \left(\frac{\kappa_i}{2} Q_i^2 + \chi_i Q_i \right) + \sum_{i,j,j<i} \frac{\kappa_{ij}^{(\text{b})}}{2} q_{ij}^2, \tag{2}$$

$$Q_i = n_i e + \sum_j q_{ij}. \tag{3}$$

In this expression, V_{short} represents the short-ranged potential (see below), V_{C} is the standard Coulomb potential, while the χ_i are the electronegativities, and the term involving κ_i is due to the atomic hardness (as in the standard QE[46]). The total atomic charges, Q_i, are the sum of an integer charge (in n_i increments of the elementary charge e) on an individual atom, as well as partial charges q_{ij} that are shared between any two bonded atoms (see Sections II A and II B). The fractional charges are antisymmetric in their indices, i.e., $q_{ij} \equiv -q_{ji}$. Single subscripts refer to quantities on individual particles (e.g., total charges or atomic properties), while double subscripts refer to quantities shared between two particles, such as a split charge, or a bond property.

The last term of Eq. (2) describes the effect of the bond hardness (as also used in the *atom-atom charge transfer* (AACT) framework[47]), and is discussed in detail below. Parameterizing the potential in terms of *both* atomic and bond properties alleviates most issues that methods only containing one or the other suffered from (see Ref.[41] and references therein, for a summary of SQE's advantages over other *electronegativity equalization methods*, which we do not repeat here).

The equations of motion are solved with a conventional velocity-Verlet algorithm, in a dedicated MD code. We use a Langevin thermostat[48] combined with stochastic damping, with a damping constant of $\gamma \Delta t = 1/600$ after the initial equilibration. The Coulomb interaction is effected in a naïve $\mathcal{O}(N^2)$ direct-sum approach, which limits the system sizes we can currently study. However, an implementation into the open-source code LAMMPS is forthcoming, which will eliminate this limitation.

For simplicity, we use the "6-12" Lennard-Jones (LJ) potential without cutoff for the short-range interactions. The electrolyte is a Kob-Andersen-like mixture,[49] in order to prevent crystallization. We use six atom types, two

for the electrolyte (positively and negatively charged), one for each electrode in its neutral state, and one for each ionized electrode species. Electrode atoms have twice the mass of solvent (electrolyte) atoms. Our choice of parameters is summarized in Table I. The unit system is explained in Section III D.

A. Bond hardness

The bond hardness $\kappa_{ij}^{(b)}$ between particles i and j, which are a distance r_{ij} apart, is parameterized by the following piecewise function:

$$\kappa_{ij}^{(b)} = \begin{cases} \kappa_{ij}^{(p)} & r_{ij} \leq r_s, \\ \kappa_{ij}^{(p)} + \kappa_{ij}^{(0)} \frac{r_l^2 (r_{ij} - r_s)^2}{r_s^2 (r_l - r_{ij})^2} & r_s < r_{ij} < r_l, \\ \infty & r_l \leq r_{ij}, \end{cases} \quad (4)$$

where r_s ($\neq 0$) and r_l are short and long cutoff radii, respectively. The symbols $\kappa_{ij}^{(p)}$ and $\kappa_{ij}^{(0)}$ denote bond parameters, constant for each bond type. In this paper the plateau value $\kappa_{ij}^{(p)} = 0$, because we are primarily concerned with metallic contacts. The functional form of the bond hardness is similar to that of Mathieu.[50] However, while we work with two variable critical radii, that work parameterized only the inner threshold with a variable scaling factor λ_{ij}. Note that Mathieu shortly simplifies $\lambda_{ij} \equiv \lambda$ as universal, i.e., not only as bond-independent but also atom-independent. The other parameter in Ref.[50] is a multiplicative factor C_{ij}. This effectively incorporates our $\kappa_{ij}^{(0)} r_l^2 / r_s^2$. Mathieu uses the van-der-Waals radius of a given atom instead of a variable outer threshold r_l.

Our parameterization of the bond hardness smoothly approaches $\kappa_{ij}^{(p)}$ at the lower threshold, while it diverges at the upper threshold, where the bond breaks. At both thresholds, the force brought about by the distance-dependence of the bond hardness has a cusp, which may lead to very small drifts in the total energy. This is discussed in detail in the previous work.[41]

B. Split charge equilibration (SQE)

Prior to calculating the Coulomb force and the MD step, in the so-called equilibration step, we update all split charges on "active" bonds for a fixed atomic configuration. A bond is classified as "active" if its bond hardness is zero or finite (but not infinite), i.e. if the bond length $r_{ij} < r_l$. Inactive bonds do not carry partial charges. We minimize the potential energy with respect to the split charge distribution by solving the homogeneous linear system of equations

$$\partial V / \partial q_{ij} = 0 \quad (5)$$

with a steepest-descent solver.[51] The potential V is that of Eq. (2). Typically, the minimization requires only

a handful of iterations. However, following an integer charge transfer (see Section II C), up to several thousand iterations can be necessary to find the split charge distribution minimizing the energy. A more efficient solver, such as a conjugate gradient method, could improve over our current implementation of the SQE.

Next, we update the total charge on each atom according to Eq. (3), based on the atomic integer charges carried, and the bond charges, and proceed with the normal MD calculation.

For the battery simulations shown below, we only allow split-charge (and integer-charge) exchange between electrode atoms. Electrolyte atoms are modeled as fixed-charge particles.

C. Integer charge transfer (ICT) in dielectric bonds

The novel feature of redoxSQE is that it allows for integer charge transfer (ICT) besides the exchange of partial charges across dielectric bonds.[41] This section briefly describes the implementation, while the subsequent section explains how we treat charge transfer across metallic bonds.

At each time step, we select all "dielectric" bonds with $r_s < r_{ij} < r_l$, i.e., any pair of atoms that is sufficiently close together to share a split charge, but not close enough to have a vanishing bond hardness ("metallic" bonds). We also exclude electrolyte atoms from participating in ICTs because they are assumed to be unreactive with the electrode for maximum battery efficiency.[2,4] Furthermore, our naïve implementation sets hard limits on the oxidation state for each atom type — in the simulations presented below we do not allow double ionization. More sophisticated rules are conceivable, but left for future work.

For each eligible bond, we draw a random number between zero and one, uniformly distributed. If it is smaller than a certain threshold (somewhat arbitrarily chosen $\equiv 1/\kappa_{ij}^{(b)}$), an ICT is attempted. This is to approximate the electron transfer rate. In more realistic simulations, this rate must be determined from quantum-chemical calculations. For a trial ICT, we increment or decrement the integer charge (i.e., the oxidation state) of each participating atom by one elementary charge, with the algebraic sign the same as the sign of the split charge between the two atoms. Then we re-equilibrate the partial charges, and calculate the system's total potential energy. If the charge transfer has lowered the energy (i.e., the system now evolves on a Landau-Zener level with lower energy), the move is accepted, otherwise it is rejected and the original state restored. A possible modification could be to accept ICTs according to some Metropolis-type condition instead.[51]

Besides changing the oxidation state, an ICT also changes the atom type. This is necessary because an ion may have different atomic properties (such as radius and interaction parameters) as well as bond characteristics

from its neutral counterpart. The type change necessitates also taking into account the short-range interaction energy for the ICT. An electrode atom only changes its type if it gains the "correct" charge. The opposite charge is absorbed by, and distributed across, metallic bonds, e.g., among connected remaining electrode atoms. For the anode, this means that an atom is stripped of a negative integer charge (i.e., one or more electrons), and becomes a cation (i.e., it changes from atom type 1 into type 3), while the negative charge remains on the anode. It serves to compensate for the positive charge accumulated by sending negative charge through the external resistor. While the net positive charge is localized on the cation, the remnant negative charge is distributed across the entire anode as split charges instantaneously, even though *formally* there is still *one particular* anode atom that carries the charge via its oxidation state, for book keeping. Conversely at the cathode, adsorbed cations receive negative integer charges transferred from the anode and are neutralized (they change from atom type 4 to type 2), as their surplus net positive charge is absorbed.

We currently re-equilibrate *all* partial charges after an ICT, which is computationally very expensive. One way to mitigate this issue could be to replace the global minimization by a local update implicating only the partial charges within some cutoff radius R^{opt}. This would mean introducing a finite signal speed for the SQE, as a split charge could be transferred a certain distance in one time step. The instantaneousness of the update would be lost, making it harder to model metals. The advantage is that it would make the operation $\mathcal{O}(N)$ instead of $\mathcal{O}(N^2)$.

In addition to the procedure described above, we draw another random number and only proceed with the trial ICT if this exceeds some threshold (for instance 0.9, to attempt an ICT only every tenth MD step), in order to alleviate a bias introduced by the order in which we query bonds. This results in two trial ICTs per atomic oscillation period, on average. In future implementations we will randomize the order of bonds for which we attempt an ICT, and fully eliminate the bias.

All in all, the maximum number of attempted ICTs per MD time step is $\sim kgNZ$, where N is the number of redox-active atoms, Z is their average coordination number, and k is the fraction of ICTs that passed picking the second random number (e.g. 10%). Lastly, $c^2 < g < c$, where $c \leq 1$ is the fraction of dielectrically bonded atoms. The factor is $g \approx c^2$ for atoms dissolved in redox-inert solvent (e.g., in the electrolyte), and $g \lesssim c$ in clusters of redox-active material.

D. Diffusion of oxidation state: integer charge transfer in metals (ICTM)

For a "metallic" bond with $\kappa_{ij}^{(\mathrm{b})} \equiv 0$, the backflow of partial charge exactly compensates the transfer of integer charge. Such a move would always be accepted because the energy is unchanged, but would still cause an expen-

sive yet unnecessary re-equilibration of split charges. In addition to the integer charge transfer across a "dielectric" bond with finite bond hardness, we therefore implement a second mode of ICT that applies to metallic bonds. We call such an operation ICTM.

We implement ICTMs such that we draw a random number for each metallic bond between two atoms otherwise eligible of an ICT. If this number exceeds a threshold, an integer charge is swapped, and immediately compensated by an equal split charge transfer in the opposite direction. Together, those transfers are energetically neutral moves in a metal. No SQE needs to be performed, and no type change occurs, so no further computations are needed. ICTMs thereby allow for "oxidation state diffusion." If the integer charge get transferred onto the front atom in the anode, we do not allow it to move away anymore.[a] As a consequence, all free negative integer charges (which could be interpreted as electrons) eventually migrate to the front atom, and are sent across the external resistor, in accordance with the real physical process. Similarly, negative integer charges emanate from the cathode's front atom and diffuse toward cations adsorbed to the electrode surface.

ICTM happens as next-neighbor hopping. It would be more meaningful if a metal cluster as a whole had an excess of integer charge (positive or negative), rather than individual atoms in a metal cluster being assigned an oxidation state. This would also reproduce realistic physics more faithfully by allowing an immediate transfer of integer charges between any two atoms connected to a metallic cluster. However, for bookkeeping and domain decomposition reasons, we stick to the current procedure, eliminating the overhead of a cluster analysis.

We emphasize that during an ICTM, only negative integer charges can diffuse through the electrodes. It is *not* possible for two initially neutral metal atoms to assume the configuration +1/-1, in contrast to the ICT in dielectrics, because one of the two atoms changes its type in such a case.

Also note that the random number is necessary to reduce (albeit not fully avoid) spurious directed motion. We perform the scan for the ICTM deterministically (e.g., atom i is always queried before atom $i + 1$), and therefore would introduce a preferred transfer order if every move was accepted.

III. SIMULATION SETUP

The configuration of a simulation with 1194 atoms at an intermediate time is qualitatively illustrated in Fig. 1.

---

[a] The condition of maximum oxidation state remains to be enforced, but is modified for the front atom to read that its *effective* oxidation state cannot be outside ± 1, i.e., the total charge after adding the "external" split charge (connected via resistor to the other electrode).

(Our default setup is somewhat smaller, and described in Section III A.) The total charge on a given atom is encoded in its color, blue being a positive charge and red a negative charge. For visual distinction, atoms are displayed at sizes that do not reflect their LJ radii. Both metallic and ionic species are visualized as much bigger than electrolyte atoms are. The latter are +1/-1 fixed-charge particles, but their charge coloring is halved, again for better contrast. The medium-sized particles are cations, while the largest particles are metallic atoms. Those may or may not carry a negative oxidation state, as negative integer charges can hop freely across metallic bonds in an ICTM (see Section II D).

The anode in this simulation initially held 5 layers of atoms; all atoms of the rightmost layer have undergone redox reactions, each donated a negative integer charge to the anode, and subsequently dissolved as a cation (roughly 45 atoms in total in this picture). The other layers are still in the middle of this process, and are only partially dissolved. The anode is positively charged, partly as a result of polarization charges as consequence of the layer of anionic electrolyte particles that wet its surface, and partly because insufficient atoms have dissolved as cations to carry away its excess charge. As expected for a metal, the charges reside primarily on the bulk surface[41].

The cathode started out with only the single fixed layer at the beginning of the simulation, but in this snapshot, about 75 pre-dissolved cations in the right half-cell have already adsorbed, donated their charge, and become part of the metallic cathode. This process is not homogeneous, the accretion occurs in clusters and produces an uneven surface with fractal-like features. If redoxSQE is properly parameterized for real materials, it can be useful in the quest to understand the detailed morphology and structure of the surface layer, as well as phenomena such as dendrite formation,[52] because it resolves the electrode-electrolyte interface. It may also be useful for intercalation studies of lithium into graphite.

Besides some minor polarization charges, the cathode is neutral overall because adsorbed cations have contributed sufficient charges to compensate for the negative charge that has traveled along the connection through the external resistor R (see Section III C).

The external load is implemented as a dedicated partial charge between an atom in the fixed layer of either electrode. This split charge between cathode and anode $q_{c\text{-}a}$ is not updated in the SQE step (see Section II B) but receives its value according to Ohm's law (if the switch is closed):

$$\dot{q}_{c\text{-}a} = U/R, \qquad (6)$$

where R is the constant external resistance, and U is the instantaneous driving voltage between the two connected front atoms. The separator only allows electrolyte particles to pass through, all others experience a repulsive force (see Section III B) if they move in between the two thin lines.
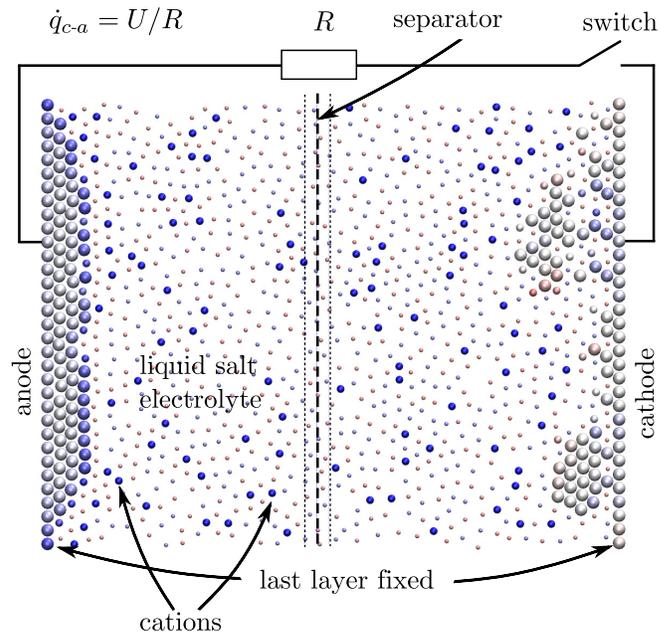
FIG. 1. Color online. Illustration of a setup with 1194 atoms. Charge is encoded in the coloring, blue meaning positive charge and red negative charge. For visual distinction, electrolyte particles are chosen smallest, independent of their LJ radii, and their charge coloring is halved. The medium-sized particles are cations, while the largest particles are metallic atoms. The separator keeps non-electrolyte atoms from moving between the two half-cells. A resistive external load R (following Ohm's law) completes the circuit.

A. Initial setup and equilibration

We briefly describe the initial configuration of our default setup, and the start of a simulation.

- We work in a two-dimensional configuration, to reduce the computational costs associated with large particle numbers. We do not apply periodic boundary conditions: our finite system is confined by fixed walls on all sides.

- Our default system contains 358 atoms overall. We have also carried out simulations with roughly double and quadruple that number.

- Of the total number of atoms, 118 are electrode atoms. Of those, half are attached to the anode and are arranged in an hexagonal lattice with three (up to five, in larger simulations) layers. Only 20 atoms are connected to the cathode in one layer, another 39 ionized cathode atoms are dissolved in the electrolyte. This represents a crude approximation to solubility equilibrium. In a historic Voltaic cell, the dissolved particles would be Cu^{2+} ions.

- The electrolyte is distributed spatially randomly in the two half cells, under the constraint that each side is initially electrically neutral.

- We choose our MD time steps such that we sample a typical thermal oscillation of an atom with ≈ 20 steps.

- Initially, we allow the electrolyte to equilibrate in each half-cell, barring all atoms from passing the separator. During this stage, the electrode particles are held stationary, and no ICTs/ICTMs are carried out.

- After $5,000$ MD time steps, the electrolyte has assumed a liquid glass-like state. We insert the separator instead of the impassable wall between the two half-cells, and allow both ICTs and ICTMs.

- We equilibrate for another $10,000$ MD time steps with all electrolyte atoms free to move, but damped at $\gamma \Delta t = 1/60$, a factor 10 more strongly than our default value.

- Finally, the electrode atoms are released, and the damping is set to its normal, low value. Only the last row of each electrode remains fixed in place. This limits the amount of charge transferable to roughly 66% of the electrode.

- We measure during the following $\approx 10^7$ MD steps.

B.   Separator

In order to isolate the two half-cells of our battery, a simple model of a *separator* is inserted. In Galvanic cells, this component (called "salt bridge" in that context) is a membrane often made of filter paper, or consists of a U-shaped glass tube filled with (possibly gelified) inert electrolyte.[4] It allows ionic species to pass through, and thus to complete the circuit with the external resistor, but prevents intermixing of dissolved electrode ions, which is often undesired.

In a present-day, commercial batteries, the two half cells are kept apart by a solid separator which is permeated by the electrolyte. In the ideal case, no electrons are allowed to pass through, but the ion conductivity is large. The separator also needs to be mechanically resistant to abuse, chemically stable in a concentrated alkaline environment (for alkaline batteries), as well as not participating in redox reactions that occur in the cell.[4]

We implement a mathematical separator such that dissolved electrode atoms and ions feel a repulsive force when they approach the barrier, but electrolyte particles are unaffected (in our model they carry charge, and can complete the circuit). If electrode particles were allowed to mix, they could exchange split charges and form salts, or even adsorb to the opposing electrode and thus create a short circuit.

Notice that the energy barrier posed by the separator is not infinitely high. As a thermally-activated process that occurs with a probability of $\exp\left(-E_{\text{sep}}/k_{\text{B}}T\right)$, individual ions are expected to still pass the barrier. We choose the separator's repulsive energy $\equiv 4$ in dimensionless units (≈ 0.6 eV). This means that an ion sitting at the separator has a probability to pass through of $\approx 4.5 \times 10^{-5}$ at our default temperature.

Our separator is a crude idealization of its real-world counterpart, even though their properties are similar. In a more realistic and properly parameterized simulation, the separator is a crucial component in its own right, and needs to be implemented carefully.

C.   External circuit

One atom of the fixed layer on each electrode is chosen as the connecting point of the "wire" to connect to the opposite electrode through an external resistor. We refer to those atoms as "front atoms." They serve as endpoints for the dedicated split charge that models the external resistor. We investigate several different modes of operation of the external circuit:

1. **switch open, no electrical connection.** The "external split charge" is constant. This mode is used for equilibration runs, as well as for aging tests of our battery.

2. **switch closed, constant resistance, discharging.** In this case, the current is determined by the instantaneous difference in chemical potential (i.e., the voltage) between the electrodes, divided by a fixed resistance. Initially, the voltage is approximately given by the difference of the electrode's electronegativities, modified by the electric field effected by the instantaneous charge distribution. If charge transfer continued in this manner indefinitely, the transferred charge would set up an opposing electric field after some time, resisting further charge transfer. However, this electric field causes a charge separation in the electrolyte, and ions rearrange to compensate it. Recall that the separator allows free exchange of electrolyte particles across the half cells. There still is a charge buildup in the electrodes, making it energetically favorable for the electrodes to shed some of that charge. This is achieved by oxidizing surface atoms on the anode, and releasing them into the solution. Analogously, at the cathode ions dissolved in the electrolyte adsorbing to the surface are reduced. This way, the electrodes are neutral again, and the voltage returns (or approaches) that of the initial state. This process ends when there are no further ions to be dissolved (and/or adsorbed).

3. **switch closed, constant resistance, charging.** This is the same setting as in the discharge case, except that we add an external voltage opposing (and overpowering) the discharge voltage. We investigate to which extent the electrodes return to

TABLE I. Normalized model parameters in our default system.[†]

atom type	χ	κ	description	base charge
1	−4	4	anode atom	0
2	4	4	cathode atom	0
3	−4	4	anode cation	+1
4	4	4	cathode cation	+1
5	1	4	electrolyte cation	+1
6	1	4	electrolyte anion	-1

bond type	$\varepsilon_{\mathrm{LJ}}$	σ_{LJ}	$\kappa_{ij}^{(0)\ddagger}$
$(1,2) - (1,2)$	1.75	1.0	3.0
$(1,2) - (3,4)$	1.0	1.0	3.0
$(1,2) - (5,6)$	1.0	1.0	
$(3,4) - (3,4)$	0.75	1.0	3.0
$(3,4) - (5,6)$	2.0	1.0	
$(5) - (5)$	0.5	1.0	
$(5) - (6)$	0.5	1.2	
$(6) - (6)$	0.5	1.0	

[†] see Section III D for details on the normalization used. [‡] The parameter $\kappa_{ij}^{(0)}$ is explained in Section II A. Electrolyte (solvent) atoms do not exchange partial charges in our model, neither among themselves nor with other species, and therefore those bonds do not have a $\kappa_{ij}^{(0)}$ assigned to them.

their previous state, and observe the battery's hysteresis. This allows to study, for instance, surface passivation.

4. **constant power or constant current.** For brevity, we only present data for discharge under constant resistance, and not under constant power or constant current, even though it is possible to model those discharge modes as well. The *constant power mode* is a good approximation to numerous real-world applications, as many electronic devices need a minimum power throughput to function properly. In the *constant current mode*, charge can continue to flow even beyond the point when the voltage drops to zero. At that point the anode and cathode reverse their roles. This setting allows studying of over-discharging behavior achievable when multiple batteries with differing remaining capacities are connected in series. In that case, the voltage of the cells that still have capacity remaining can drive the empty ones into pole reversal.

D. Unit system and parameters

Throughout this work we use dimensionless parameters. In order to facilitate the interpretation of the presented data, this section provides ballpark estimates for representative values of real materials.

The unit of charge can be associated with the elementary charge $[Q] = 1.6 \times 10^{-19}$ C. For the unit of length, we choose $[l] = 2.3$ Å. This is to approximate the Lennard-Jones parameter σ_{LJ} used for Cu-Cu interactions in the literature.[53] The value for Zn is comparable

in magnitude albeit slightly larger. We define the unit of mass as the atomic mass of copper, $[m] \approx 64$ amu. As last independent unit, the energy is normalized to $[E] = 0.16$ eV ≈ 15.5 kJ/mol. Again, this is so that the interaction between metallic electrode particles is comparable to values for ϵ_{LJ}(Cu-Cu) in the literature.

Together, length, charge, energy, and mass specify a complete set of units for our purposes. Derived units are the unit of current $[I] \approx 350$ nA, $[R] \approx 3.3 \times 10^5$ Ω as resistance normalization, $[U] \approx 160$ mV as unit of voltage, and $[t] \approx 0.5$ ps as unit of time. Our battery demonstrator operates at a temperature $T \approx 740$ K, comparable to a liquid-salt battery.[54]

With these choices of units, the default value for the electronegativity difference between anode and cathode is $\Delta\chi = 1.28$ V, which is close to many standard cells, e.g., alkaline (1.5 V) or NiMH (1.2 V) batteries. In the default parameterization, we use atomic hardnesses of $\kappa = 0.64$ eV, which is much smaller than typical values, e.g., $\kappa_{\mathrm{Cu}} = 7.3$ V/e and $\kappa_{\mathrm{Li}} = 4.7$ V/e.[55] Moreover, unlike real systems, our standard values for χ and κ do not reproduce a neutral dissociation limit of dimers, because $\kappa/2 - \chi > 0$, as can be seen from Eq. (2). We made this choice of parameters to accelerate the generation of ions. At the same time, we ensured in selected simulations that the qualitative features of the discharge curves remained unchanged for much larger values of κ and smaller values for χ (see Section IV).

In order to compare our results with macroscopic systems, our model battery would have to be scaled up by a factor of roughly 10^{21}. In principle, each spatial dimension can have a different scaling factor, however, for simplicity, one may assume the same factor of 10^7 in each direction. An inherent ambiguity of how to scale the direction normal to the interfaces remains. One way to scale the simulation is to take our nano-battery as an electrical element and connect 10^7 of them in series, and 10^{14} in parallel. This method yields an overall open-circuit voltage (OCV) of 12.8 MV for our default choice of parameters ($\Delta\chi = 8$). Connecting the external resistors in a similar fashion as the batteries leads to a scaled resistance of 66 Ω (microscopic system: $R = 2000$ in dimensionless units). Consequently, the nominal macroscopic discharge current would be 70 kA. Discharge now takes 10^7 times longer (about 0.1 s) than in the microscopic case, because the total number of transferrable charges increases by 10^{21}, while the current only increases by 10^{14}.

Alternatively, one can scale both battery and resistor *as a whole* in each spatial dimension. This way one retains the value for the macroscopic resistance of 66 Ω. In contrast, the voltage in this case still has its microscopic value of 1.28 V. The resulting nominal macroscopic discharge current is 7 mA, which is not very much lower than real-world currents. Now the discharge would take 10^{14} times longer than in our microscopic model (about 14 d). We stress that the discharge characteristics of such a scaled-up version of our model battery would be

different from the ones presented in this work because the electrode surface-to-volume ratio would be much reduced in the macroscopic battery, among other reasons.

E. Limitations and code efficiency

At this stage, our code does not contain very sophisticated physics, besides some rudimentary approximation to chemistry, and the ability to model redox reactions. We only consider two-body forces, and leave dihedral and torsional interactions for future work. Moreover, at short distances we do not screen the Coulomb interaction, even though the wave functions of atoms overlap in such a situation, and the point-charge approximation breaks down.

Our current implementation is rather primitive, and can be made to run much more efficiently. Particular points to note are the following. We do not cut off the Lennard-Jones interaction, and treat it as a long-range force. We also compute the Coulomb interactions with an expensive $\mathcal{O}(N^2)$ direct-sum algorithm, and could save computing time with a more efficient approach. However, at this developmental stage, an ICT necessitates a complete recomputation of all split charges in the system. The solution of the large linear system is more expensive than the setup of the system matrix, for which the Coulomb term is calculated. Two ways to make this cheaper come to mind: the first is to use a more efficient algorithm for the solution of the linear system.[51] Second, we could update only the split charge in vicinity of the ICT, instead all of them. The resulting error can be estimated and is bounded; the change in the split charges, and thus the error as a result of restricting the update distance, drops off exponentially with increasing cutoff distance.[56]

However, all of those optimizations are left for future work; even without them (albeit for small systems) the method yields encouraging results.

IV. RESULTS

In this section we describe the results gleaned from our simulations. We emphasize again that our model is more a proof-of-concept rather than a faithful representation of a real battery. As such, the parameters have not been chosen appropriate for any specific material. Our intent is to demonstrate that redoxSQE nevertheless reproduces generic properties of macroscopic batteries without further input.

The model necessarily has a number of parameters, albeit not nearly as many as some mesoscopic porous electrode simulations.[57] Some are microphysical and chemical parameters, for instance the LJ properties of the materials, the atomic hardness, or the electronegativities. Those are in principle readily parameterizable or even measurable quantities for real materials, but we elect to use representative values (see Section III D), rather

than to model specific materials. The parameters associated with the bond hardness are not as easily measured, but can in principle be fit to values found from measurements and *ab initio* quantum-chemical DFT simulations, such as ESP partial charges, Hirshfeld-I charges, and dipole moments of molecules.[42,50,58] Other parameters are implementation choices such as number of atoms to model, the electrode setup, and the geometry of our cell. Finally, we vary parameters that have a documented and experimentally accessible impact on battery performance, such as the temperature, and whether the battery is discharged continuously or in pulses. The dependence on this last set of parameters results naturally and self-consistently with our method, and does not have to be put in implicitly or explicitly.

A. Dependence on internal model parameters

In this section, we explore the dependence of the discharge characteristics on internal model parameters, while Section IV B focuses on external parameters. In SQE, the difference in electronegativity between atoms of two metals determines the open-circuit voltage (OCV). In all following plots, the voltage is normalized to this value.

Separating a diatomic molecule adiabatically results in neutral products for each pair of stable elements. However, if $\Delta\chi > (\kappa_1 + \kappa_2)/2$, it is energetically favorable that negative integer charge (i.e., an electron) remains on the more electronegative partner.[39] For a multi-atom system, the expression for the neutral dissociation limit is not as simple anymore, because the atomic hardness is reduced in an ensemble.[39] In our default system, we choose $\chi = \pm 4$ and $\kappa = 4$, in order to facilitate the formation of ions at the anode, even though these values lead to a violation of the neutral diatomic dissociation limit. However, Fig. 2 shows that this does not have a large impact on the discharge curve: even for $\chi = \pm 1.9$ and $\kappa = 4$, satisfying the neutral dissociation limit, the qualitative picture remains.

We note that the large scatter stems from the small number of atoms in our simulations. If we dissolve 50% of the total number of anode atoms, we have 30 cations in solution, and a statistical uncertainty of $\mathcal{O}\left(1/\sqrt{30}\right)$, i.e., $\approx 15\%$. The stochastic error reduces by a factor of two by averaging over 4 independent realizations.

In Section III D we reported representative values as normalization for the dimensionless parameters used in this work, and noted that our standard value for the atomic hardness κ was rather low compared with values for real materials. In Fig. 3, we show the discharge behavior of our nano-battery for $\kappa = 10$ and $\Delta\chi = 8$, again satisfying the neutral diatomic dissociation limit. In this case, ion formation (i.e., redox reactions at the electrode surface) takes much longer, and the external resistor needs to be scaled up in order to get similar behavior. A factor of 2.5 in κ is approximately compensated
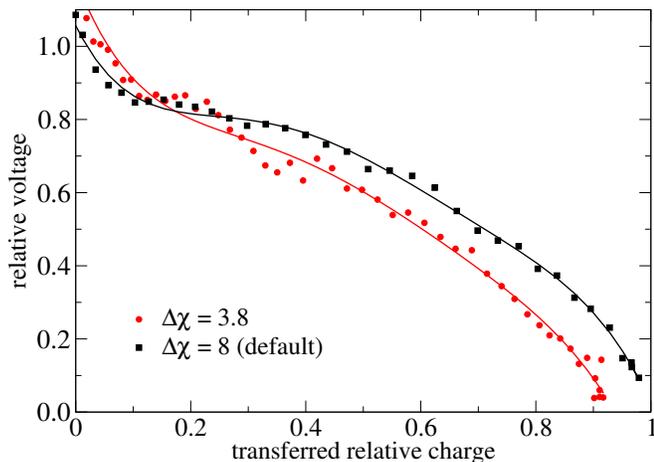
FIG. 2. Color online. Discharge curves of our battery demonstrator for different electronegativities. The data represents an average over several independent runs, and each point is averaged over many MD time steps. The solid lines are inserted to guide the eye. Our default model has $\chi = \pm 4$ and $\kappa = 4$ which violates the dissociation limit: dissociating a diatomic molecule will produce ionic instead of neutral products. The runs with $\chi = \pm 1.9$ do not have this flaw, and still exhibit similar discharge behavior, even though the likelihood of anode atoms to dissolve as ions is reduced. The scatter is larger for smaller χ since the theoretical voltage is reduced by a factor of 2.1, and thus the signal-to-noise ratio halves, even though in this case we average over 8 different realizations.

FIG. 3. Color online. Discharge curves of our battery demonstrator varying the atomic hardness. The data represents an average over 4 independent runs, and each point is averaged over many MD time steps. The solid lines are inserted to guide the eye. Runs with $\kappa = 10$ (approximating real atomic hardnesses more closely than our standard value) require much higher external resistances in order to achieve similar discharge behavior as for the default value of κ. For larger values, the plateau is much less pronounced. The $\kappa = 10$ runs are scaled down by a factor of 1.3 so that they overlap the $\kappa = 4$ curves better (see text).

by a factor of 25 in resistance. In addition, the plateau is far less pronounced for the larger κ, and the discharge proceeds faster. Note that the initial voltage in this case is also higher. The reason is that some pre-dissolved cations adsorb to the electrode immediately, and cause a greater difference in chemical potential, and therefore OCV. In order to have the curves overlap for better visual comparability, we scaled down the results for $\kappa = 10$ by a factor of 1.3, which compensates for the larger OCV.

In Fig. 4, we exemplify the influence of the LJ parameters on the discharge curve. We vary $\varepsilon_{\mathrm{LJ}}$ between electrode ions and the electrolyte. A higher value means that it is favorable for an ion to surround itself with electrolyte atoms, as opposed to other atoms with which its $\varepsilon_{\mathrm{LJ}}$ is lower. The discharge curve has a higher and more extended plateau for a smaller $\varepsilon_{\mathrm{LJ}}$, caused by the reaction on the cathode, where a lower value means that it is more likely that an ion is adsorbed to the electrode. At the anode, the opposite should be the case. There, a *greater* $\varepsilon_{\mathrm{LJ}}$ should make it more likely for an ion to be dissolved into the electrolyte. However, the test case with $\varepsilon_{\mathrm{LJ}}^{\mathrm{anode}} > \varepsilon_{\mathrm{LJ}}^{\mathrm{cathode}}$ changes the plateau only marginally. We conclude that the cathode reaction is more important in this respect. If the method in implemented into more sophisticated software (e.g., LAMMPS), it can be used with more realistic force fields for particle-particle interactions than what is used herein (i.e., simple two-particle LJ interaction).

FIG. 4. Color online. Discharge curves of our battery demonstrator varying the LJ parameter $\varepsilon_{\mathrm{LJ}}$ between electrode ions and the solvent (electrolyte). The data represents an average over 4 independent runs, and each point is averaged over many MD time steps. The solid lines are inserted to guide the eye. This figure exemplifies the optimization potential that lies in picking optimal materials for the battery: a smaller affinity between electrolyte and cations boosts battery performance significantly. The anode half-cell reaction seems less important in this respect, as optimizing it does not yield a comparable additional improvement.

While the LJ parameters in principle can be deduced from experiments, pinning down the parameterization details of the bond hardness so that they match experiments or *ab initio* results is harder to accomplish.[42,50,58] Fortunately, those parameters do not affect the results very strongly, as evidenced in Fig. 5. A change by 50% in the cutoffs for $\kappa_{ij}^{(b)}$, given in Eq. (4), does not make a big difference: all curves nearly overlap with those of our default model in the practically relevant regime (until a discharge of $\gtrsim 60\%$). In Fig. 6 we scale the parameter $\kappa_{ij}^{(0)}$ up and down by a factor of 3. The results do not depend sensitively on this choice, either.[b] These results indicate that the detailed form of $\kappa_{ij}^{(b)}$ is not of great importance. The only criterion is that the next-nearest neighbor should not be connected with a dielectric bond. This translates into the "long" cutoff r_l be substantially smaller than the distance to the next-nearest neighbors. Otherwise a great number of additional ICTs will be attempted. This does not change the result, either, but the computations will be slowed down tremendously.
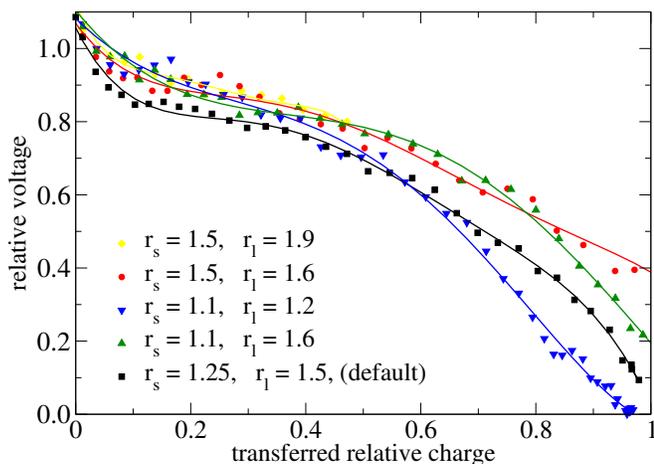
FIG. 5. Color online. Discharge curves of our battery demonstrator for various cutoffs in the parameterizations of the bond hardness. The data represents an average over 4 independent runs, and each point is averaged over many MD time steps. The solid lines are inserted to guide the eye. A change in the cutoffs entering $\kappa_{ij}^{(b)}$ by 50% does not significantly alter the results: all curves resemble our default model. This indicates that the detailed form of $\kappa_{ij}^{(b)}$ is not of great importance.

We tested the effect of various other simulation parameters. For brevity we only describe the results without including additional figures.

- We increased the damping by a factor of 10 without observing any qualitative nor quantitative changes in the discharge behavior. This means that the

---

[b] Note that the cutoffs and $\kappa_{ij}^{(0)}$ are not fully independent. As shown in Eq. (4), all three have a scaling effect on $\kappa_{ij}^{(b)}$.
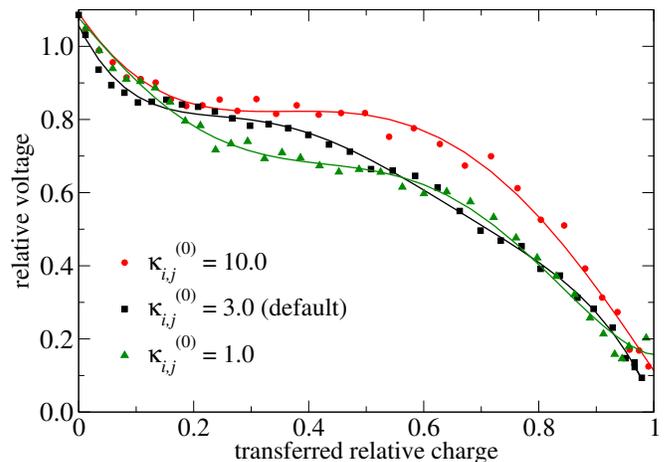
FIG. 6. Color online. Discharge curves of our battery demonstrator for different parameterizations of $\kappa_{ij}^{(b)}$. The data represents an average over 4 independent runs, and each point is averaged over many MD time steps. The solid lines are inserted to guide the eye. We vary the parameter $\kappa_{ij}^{(0)}$ by a factor of 3 to greater and smaller values. The results do not depend sensitively on this choice.

Langevin damping only has the desired effect to limit the battery heating up as a consequence of energy release, but is not strong enough to influence the dynamics much.

- A variation in the LJ parameters between electrode metals and their ions by 50% also did not alter the characteristics significantly. Our choice stems from the idea that ions are bound less tight to the electrodes than neutral metal atoms. This again aids the release of cations into solution.

- We modified the relative importance of the different contributions to the potential energy. A decrease of the effect of the Coulomb energy by up to 33% (and thereby a corresponding increase in relative importance of the other effects) did not influence the discharge curve qualitatively.

- Similarly, scaling the battery in one direction, doubling the amount of electrolyte without adding more electrode atoms, nor changing the electrode surface area modified the results significantly. We conclude that we are not hampered by a lack of electrolyte. However, reducing the number of electrolyte atoms to half its default value will limit the number of ions that can be dissolved, and deteriorate our battery demonstrator's performance.

B. Dependence on external factors, rates and aging

An ideal battery should retain its theoretical voltage until the active material has been used up, that is, until the anode is completely dissolved, or all free cations ions

have been reduced and adsorbed at the cathode surface. Only at this point should the voltage drop to zero.[4] In reality, batteries have an internal resistance, and both the electrolyte and the electrode are polarizable. The former reduces the actual voltage drop across the battery, while the latter is responsible for forming Helmholtz double layers at the electrode surfaces, and thus depleting some of the battery's capacity.[c] Such effects cause both the actual working voltage as well as the usable capacity to be reduced from their theoretical limits. Additionally, a battery's voltage also depends on the discharge current, such that a higher discharge current will decrease the discharge voltage. In Fig. 7 we show curves demonstrating this behavior. At high external resistance (i.e., low discharge current), the voltage stabilizes at $\lesssim 90\%$ of its theoretical voltage until about 70% of the available charge has been transferred, at which point it drops quickly. At medium resistance (our default model), the voltage does not feature such a pronounced plateau, but still has $\approx 60\%$ of its voltage at 60% discharge. In contrast, the voltage for a discharge at high currents decreases much more steeply. These curves are reminiscent of those presented in Refs.[4,10,14,17], the discharge curves of Panasonic's zinc carbon batteries,[59] and those of Duracell's alkaline batteries.[60] If the internal resistance of

[c] A battery's capacity (in Wh) is the integral under the curve voltage vs. transferred charge, the discharge curve.
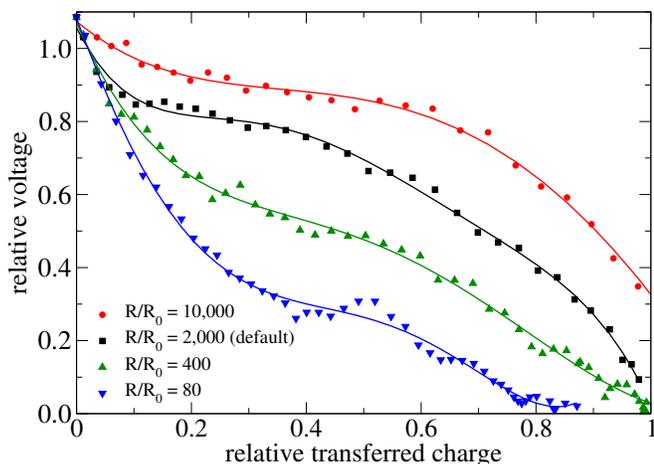
FIG. 7. Color online. Discharge curves of our battery demonstrator with different external resistors. The data represents an average over 4 independent runs, and each point is averaged over many MD time steps. The solid lines are inserted to guide the eye. The higher the external resistance, the closer the battery's behavior approaches the theoretical discharge curve. In this property, and the shape of the discharge curve, our nano-battery resembles a macroscopic battery. Initially, the voltage declines sharply, as the electrodes are charged before ions are dissolved or adsorbed, respectively. This is followed by an extended plateau when the voltage stays constant as the charge transfer through the external resistance is balanced by an equal amount of ion transfer in the electrolyte. At the same time, additional charge on the electrodes is now compensated by dissolving and adsorbing ions. Finally, another steep decline concludes the discharge, as the electrodes are consumed and their surfaces passivated.

a battery exceeds the external resistance, we effectively have a short circuit, and the battery will discharge as a capacitor, with an initial exponential decay of the voltage.

A factor that is of crucial importance to real batteries is the temperature at which they operate. Electric vehicles need to be able to reliably operate at temperatures ranging from $\lesssim -30°$ C all the way to $\gtrsim 50°$ C. However, low temperatures decrease both the actual voltage as well as the battery's capacity.[4] Figure 8 shows the temperature dependence of our battery demonstrator. As in macroscopic batteries, the voltage is closer to the theoretical voltage for high temperatures, while it is significantly reduced for lower temperatures. The battery's capacity — the area under the curve — decreases by $\approx 42\%$ if the temperature is lowered by 25%, and increases by $\approx 46\%$ if the temperature is raised by 50% from our default temperature. We note that our temperature of $T = 740$ K is quite large, comparable to liquid-salt batteries,[54] but greatly exceeds room-temperature. However, we once more emphasize the qualitative nature of our findings, not their realism.

In Fig. 9 we present aging studies of our model battery. It is common in some electronic devices to intersperse recuperation periods with discharge periods. During this time, polarization effects are reduced and some of the initial voltage can be recovered.[4] In order to examine the effects of intermittent discharge on our model battery, we discharge it until a certain amount of current has been drained, and then open the switch and let the system age. Were it an electrochemical capacitor, no recovery of voltage would be expected. However, we see a nearly full recuperation of the initial voltage (with some fluctuations). This voltage is held for $\approx 10^7$ time steps. After some time, ions manage to pass through the separator (which is a thermally activated process, see Section III B). This causes a voltage drop that contributes to self-discharge. In a macroscopic battery this process will not take place as quickly as in our nanoscale device, but one mechanism for is elucidated in our model. The use of redoxSQE furthermore allows us to study the morphology changes that
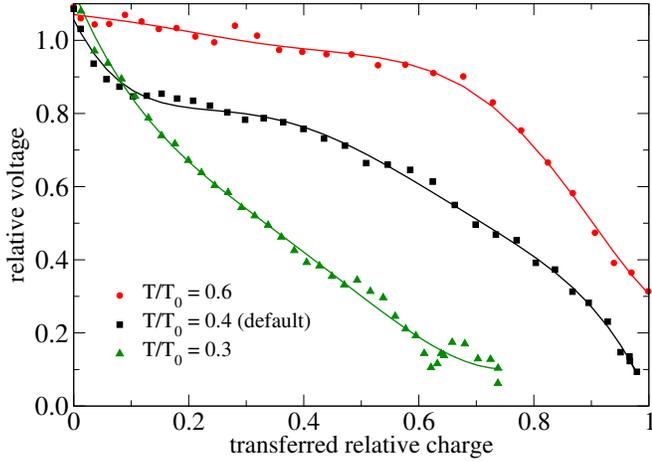
FIG. 8. Color online. Discharge curves for different system temperatures. The data represents an average over 4 independent runs, and each point is averaged over many MD time steps. The solid lines are inserted to guide the eye. At lower temperatures ($T \approx 550$ K), the battery's performance deteriorates dramatically, while it improves at a higher temperature ($T \approx 1,100$ K).

the electrodes have undergone, how much surface material is passivated and other microphysical parameters of interest.
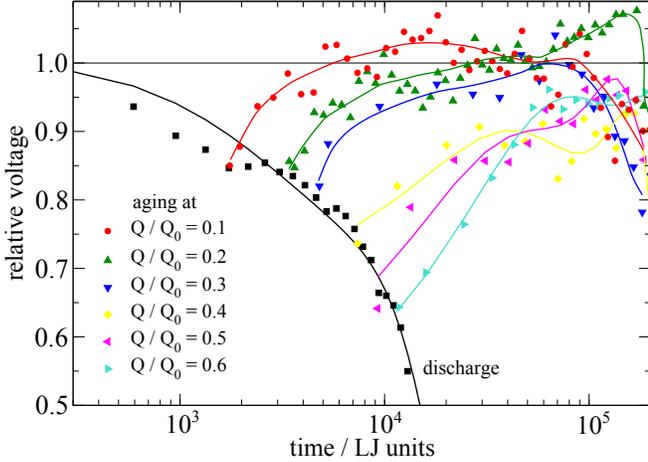
FIG. 9. Color online. Aging studies of our model battery. The data represents an average over 4 independent runs, and each point is averaged over many MD time steps. The solid lines are inserted to guide the eye. After a certain amount of charge has flown through the resistor, the switch is opened and the relaxation observed. In all cases, the voltage recuperates to almost the nominal value (with fluctuations). This behavior would not be present in a capacitor, which does not recover its charge after the circuit is broken. After a certain time, ions pass the separator and cause self-discharge, which gives rise to a voltage drop.

All reactions occurring in our system are micro-reversible, therefore our nano-battery is a secondary, rechargeable cell. Figure 10 shows what happens if

we recharge our battery demonstrator. We discharge our default system at the default constant resistance ($R = 2000$). After $\approx 60\%$ of the total capacity (35 out of 59 integer charges) have been transferred through the external circuit, a charging current is switched on, in opposite direction as the discharge current. We choose its magnitude approximately three times the "average" discharge current (averaged over the complete discharge of the same model).

We consider two cases: in the first, the charging current is switched off again when the voltage reaches ≈ 1.1 times the OCV. In this case, the discharge curve in the second cycle has deteriorated compared with the initial discharge; the battery has degraded. The reason is that the electrodes do not fully return to their initial state during the charging, but merely the electrolyte reconfigures to balance the dissolved cations in either half cell. If one were to let the system relax after charging, some dissolved cations would return to the anode, and some material deposited on the cathode would also dissolve again. A realistic all-atom simulation using redoxSQE can be useful in the investigation of dendrite formation in this process which can short-out and destroy Li-ion batteries.[1,2,61] It could also help to study the cycling behavior of batteries, and their degradation.

In the second case, all charge is transferred back. Then, the voltage is much higher than the initial OCV. This behavior is also seen in real batteries,[4] and stems from a buildup of a polarization layer opposite that which forms during discharge. In this case, the second discharge is not very dissimilar from the first one.
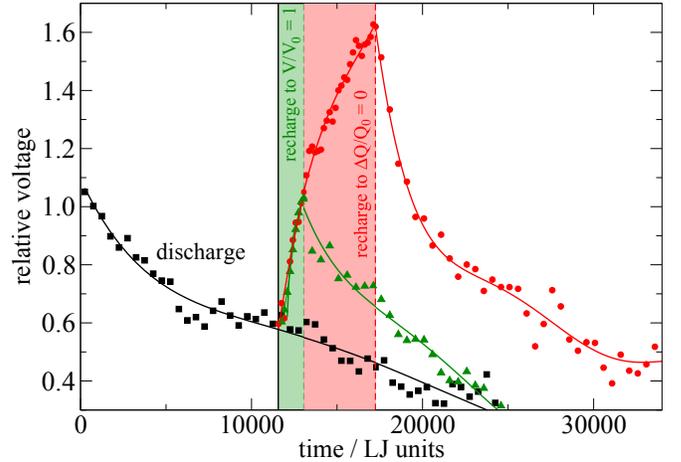
FIG. 10. Color online. Discharge and recharge curves of our battery demonstrator. The lines are inserted to guide the eye. Our default system is discharged (black squares) at constant resistance, until $\approx 60\%$ of the total charge has been transferred through the external resistor. At this point, a constant external charging current is switched on; its magnitude is chosen approximately three times the (average) discharge current. This charging current is switched off again when the voltage reaches the OCV (middle, green triangles), or when all charge has been transferred back (top, red circles). In the latter case, the voltage is much higher than the initial OCV.

V. DISCUSSION AND CONCLUSIONS

In this work, we demonstrated that the redoxSQE method[39] can model redox reactions in an atomistic molecular dynamics setting, and used it to simulate a nanoscale battery demonstrator. Even though we did not use a parameterization describing any real energy materials, we reproduced generic discharge curves of macroscopic batteries. For example, lower operating temperatures reduce the effective capacity of a battery. Higher discharge rates have the same effect, but the voltage recuperates when the battery is aged (e.g., discharged in pulses). Upon recharging, the battery performance degrades slightly, and the electrode surface morphology changes during the battery's operation.

Some internal model parameters are not fully accessible experimentally, such as the bond hardness, which is an ad-hoc parameter arrived at in a top-down fashion, by describing the bond-breaking behavior. In a theoretical work, Verstraelen et al.[45] connect the bond hardness to parameters computed with atom-condensed DFT, i.e., derive it in a bottom-up fashion, and may help to motivate this parameter and ascertain its value quantitatively. We showed that, for the battery demonstrator, the results neither depend strongly on the detailed implementation, nor on the precise value of the bond hardness.

The *atomic* hardness plays an important role in determining the time scales of ion formation at the electrode interface, and thus determines partly the internal resistance of the battery, while the electronegativity difference sets the open-circuit voltage. But changing those two parameters does not alter the qualitative discharge picture. Rather, the battery's behavior is predominantly defined by external quantities such as temperature, rates of discharge. We obtain results similar to the intermittent discharge mode,[4] and see self-discharge when the battery is aged too long.

In a recent previous paper we applied the same technique to case studies in contact electrification between two clusters of ideal metals and ideal dielectrics, respectively, showcasing its ability to simulate history-dependence.[41] RedoxSQE reproduces charge hysteresis effects during approach and retraction.

One shortcoming of our current implementation is that the electrolyte is modeled with fixed-charge particles that do not participate in split charge exchange, nor in ICTs. It is an ideal insulator for electrons, and the lack of electronic conduction leaves only penetration of the separator by ions as self-discharge mechanism. This idealization will be abandoned in future work.

Further development effort will need to be expended on optimizing the method (see Section III A, and Ref.[39]) to make multi-million atom simulations possible. An implementation into LAMMPS is planned. In order to simulate specific materials or battery setups (such as alkaline batteries, or Li-ion rechargeables), much chemically-specific parameterization will need to be done.[40,42,50,58] Furthermore, more realistic empirical many-body force fields are necessary for realistic all-atom simulations. We point out that the model in its current implementation can best describe non-directed interactions, as they are prevalent for instance in Alkali batteries, with their isotropic reactions of s-orbitals.

Notwithstanding those necessary improvements, it is encouraging that the method already reproduces generic features of batteries. Mesoscale battery models require many assumptions and intimate knowledge of the materials in question, and cannot answer fundamental microphysical questions. DFT/MD methods, on the other hand, have been used for highly detailed and isolated problem aspects, but need to stay away from the electrode-electrolyte interface where redox reactions take place. Arguably, this is the most interesting region, as it determines not only the ultimate cell performance, but also is where cell degradation takes place. Harris et al.[19] write "*the ability to predict cell degradation remains a challenge because so many unaccounted for and seemingly unrelated micro-scale degradation mechanisms have been identified or postulated. [...] Without ... theoretical analysis, cause-and-effect relationships between observation and degradation pathway can be difficult to demonstrate.*" RedoxSQE is a first step toward filling this gap. It allows to model all aspects of a (microscopic) battery in one simulation, and gather insights into the processes happening at the electrode-electrolyte interface.

Besides modeling an entire all-atom battery, redoxSQE could also serve as part of hybrid multiscale schemes, where bulk phenomena inside an electrode or within the electrolyte are computed with a mesoscopic model while the electrochemical activity is tackled by redoxSQE.

Daniel's Handbook of Battery Materials[62] states as requirement for a generic life estimation model that it "*must relate the measured cell performance at any given time to a combination of ... effects.*" In conjunction with a more realistic force field and with a proper parameterization of the materials, redoxSQE holds promise to enable the study of battery degradation and the optimization of battery performance.

[1] J. Tarascon and M. Armand, Nat. **414**, 359 (2001).
[2] M. Armand and J. M. Tarascon, Nat. **451**, 652 (2008).
[3] R. A. Huggins, *Advanced Batteries* (Springer, New York, 2009).
[4] D. Linden and T. B. Reddy, *Linden's Handbook of Batteries* (McGraw-Hill Professional, New York, 2010).
[5] B. Scrosati and J. Garche, J. Power Sources **195**, 2419 (2010).
[6] J. Newman and C. Tobias, J. Electrochem. Soc. **109**, 1183 (1962).
[7] R. Pollard and J. Newman, J. Electrochem. Soc. **128**, 491 (1981).
[8] M. Doyle, T. Fuller, and J. Newman, J. Electrochem. Soc. **140**, 1526 (1993).
[9] T. Fuller, M. Doyle, and J. Newman, J. Electrochem. Soc. **141**, 1 (1994).

[10]M. Doyle, J. Newman, A. Gozdz, C. Schmutz, and J. Tarascon, J. Electrochem. Soc. **143**, 1890 (1996).

[11]R. Darling and J. Newman, J. Electrochem. Soc. **145**, 990 (1998).

[12]G. Botte, V. Subramanian, and R. White, Electrochim. Acta **45**, 2595 (2000).

[13]P. Ramadass, B. Haran, P. Gomadam, R. White, and B. Popov, J. Electrochem. Soc. **151**, A196 (2004).

[14]V. Srinivasan and J. Newman, J. Electrochem. Soc. **151**, A1517 (2004).

[15]S. Santhanagopalan, Q. Guo, P. Ramadass, and R. White, J. Power Sources **156**, 620 (2006).

[16]C. Wang, U. S. Kasavajjula, and P. E. Arce, J. Phys. Chem. C **111**, 16656 (2007).

[17]C.-W. Wang and A. M. Sastry, J. Electrochem. Soc. **154**, A1035 (2007).

[18]X. Zhang, W. Shyy, and A. M. Sastry, J. Electrochem. Soc. **154**, A910 (2007).

[19]S. J. Harris, A. Timmons, D. R. Baker, and C. Monroe, Chem. Phys. Lett. **485**, 265 (2010).

[20]I. V. Thorat, D. E. Stephenson, N. A. Zacharias, K. Zaghib, J. N. Harb, and D. R. Wheeler, J. Power Sources **188**, 592 (2009).

[21]A. Awarke, M. Wittler, S. Pischinger, and J. Bockstette, J. Electrochem. Soc. **159**, A798 (2012).

[22]I. Abou Hamad, M. A. Novotny, D. O. Wipf, and P. A. Rikvold, Phys. Chem. Chem. Phys. **12**, 2740 (2010).

[23]D. Liu and G. Cao, Energy & Environmental Science **3**, 1218 (2010).

[24]C. A. J. Fisher, V. M. H. Prieto, and M. S. Islam, Chem. Mat. **20**, 5907 (2008).

[25]B. L. Ellis, K. T. Lee, and L. F. Nazar, Chem. Mater. **22**, 691 (2010).

[26]M. S. Islam, Phil. Trans. R. Soc. A **368**, 3255 (2010).

[27]G. Hautier, A. Jain, S. P. Ong, B. Kang, C. Moore, R. Doe, and G. Ceder, Chem. Mater. **23**, 3495 (2011).

[28]R. Tripathi, G. R. Gardiner, M. S. Islam, and L. F. Nazar, Chem. Mat. **23**, 2278 (2011).

[29]S. Lee and S. S. Park, J. Phys. Chem. C **116**, 6484 (2012).

[30]C. Ouyang, S. Shi, Z. Wang, X. Huang, and L. Chen, Phys. Rev. B **69** (2004).

[31]C. Arrouvel, S. C. Parker, and M. S. Islam, Chem. Mat. **21**, 4778 (2009).

[32]W. J. Mortier, K. van Genechten, and J. Gasteiger, J. Am. Chem. Soc. **107**, 829 (1985).

[33]V. May and O. Kühn, *Charge and Energy Transfer Dynamics in Molecular Systems* (Wiley-VCH, Weinheim, 2011).

[34]E. Runge and E. K. U. Gross, Phys. Rev. Lett. **52**, 997 (1984).

[35]E. Gross and W. Kohn, Adv. Quant. Chem. **21**, 255 (1990).

[36]R. Bauernschmitt and R. Ahlrichs, Chem. Phys. Lett. **256**, 454 (1996).

[37]J. Blumberger, Y. Tateyama, and M. Sprik, Comput. Phys. Commun. **169**, 256 (2005).

[38]S.-C. Mao, J.-Q. Qu, and K.-C. Zheng, Chin. J. Chem. Phys. **25**, 161 (2012).

[39]M. H. Müser, Eur. Phys. J. B **85** (2012).

[40]T. Verstraelen, E. Pauwels, F. De Proft, V. Van Speybroeck, P. Geerlings, and M. Waroquier, J. Chem. Theory Comput. **8**, 661 (2012).

[41]W. B. Dapp and M. H. Müser, accepted by Eur. Phys. J. B (2013).

[42]R. A. Nistor, J. G. Polihronov, M. H. Müser, and N. J. Mosey, J. Chem. Phys. **125** (2006).

[43]P. T. Mikulski, M. T. Knippenberg, and J. A. Harrison, J. Chem. Phys. **131** (2009).

[44]M. T. Knippenberg, P. T. Mikulski, K. E. Ryan, S. J. Stuart, G. Gao, and J. A. Harrison, J. Chem. Phys. **136** (2012).

[45]T. Verstraelen, P. W. Ayers, V. Van Speybroeck, and M. Waroquier, J. Chem. Phys. **138**, 074108 (2013).

[46]A. K. Rappé and W. A. Goddard, J. Chem. Phys. **95**, 3358 (1991).

[47]R. Chelli, P. Procacci, R. Righini, and S. Califano, J. Chem. Phys. **111**, 8569 (1999).

[48]T. Schneider and E. Stoll, Phys. Rev. B **17**, 1302 (1978).

[49]W. Kob and H. Andersen, Phys. Rev. E **51**, 4626 (1995).

[50]D. Mathieu, J. Chem. Phys. **127** (2007).

[51]W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes* (Cambridge University Press, New York, 2007).

[52]I. Valov, R. Waser, J. R. Jameson, and M. N. Kozicki, Nanotechnol. **22** (2011).

[53]H. Hwang, O. Kwon, and J. Kang, Solid State Commun. **129**, 687 (2004).

[54]D. J. Bradwell, H. Kim, A. H. C. Sirk, and D. R. Sadoway, J. Am. Chem. Soc. **134**, 1895 (2012).

[55]D. C. Ghosh and N. Islam, Int. J. Quantum Chem. **110**, 1206 (2010).

[56]R. A. Nistor and M. H. Müser, Phys. Rev. B **79** (2009).

[57]T. Menzel, D. Willkomm, and A. Wolisz, in *Proc. of the CONET 2011 workshop in conjunction with CPSWeek 2011* (CPSWeek, Chicago, USA, 2011).

[58]T. Verstraelen, V. Van Speybroeck, and M. Waroquier, J. Chem. Phys. **131** (2009).

[59]Panasonic Corporation, "Panasonic zinc carbon batteries," (2009), online; accessed on Nov 19, 2012.

[60]Procter & Gamble, "Duracell alkaline-manganese dioxide battery technical bulletin," (2012), online; accessed on Nov 19, 2012.

[61]R. Bhattacharyya, B. Key, H. Chen, A. S. Best, A. F. Hollenkamp, and C. P. Grey, Nat. Mater. **9**, 504 (2010).

[62]C. Daniel and J. O. Besenhard, *Handbook of Battery Materials* (Wiley-VCH, Weinheim, 2011).