

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Drinking Water in the Rural Developing World: Improving Ion Flux in Capacitive Deionization and Understanding Technical Interventions

Permalink

<https://escholarship.org/uc/item/2vq9p6s3>

Author

Rausch, Adam John

Publication Date

2017

Peer reviewed|Thesis/dissertation

**Drinking Water in the Rural Developing World: Improving Ion Flux in
Capacitive Deionization and Understanding Robust Technical Interventions**

by

Adam John Rausch

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering-Civil and Environmental Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Ashok Gadgil, Chair

Associate Professor Isha Ray

Professor David Sedlak

Summer 2017

**Drinking Water in the Rural Developing World: Improving Ion Flux in
Capacitive Deionization and Understanding Robust Technical Interventions**

Copyright 2017
by
Adam John Rausch

Abstract

Drinking Water in the Rural Developing World: Improving Ion Flux in Capacitive Deionization and Understanding Robust Technical Interventions

by

Adam John Rausch

Doctor of Philosophy in Engineering-Civil and Environmental Engineering

University of California, Berkeley

Professor Ashok Gadgil, Chair

Today, a billion people lack access to safe drinking water. Climate change, a host of unsustainable practices, and emerging contaminants threaten to exacerbate the problem. For the developing-world communities disproportionately affected, limited infrastructure and resources often preclude the developed-world solutions. New technologies will be necessary. Unfortunately, even mature technologies often fail in these settings. This work presents two efforts: one to advance understanding of robust technical interventions and one to improve water treatment technology for such settings.

In the Poshiir River watershed study, we built comparative case studies of single-village piped-water drinking water schemes in rural Maharashtra. In-depth datasets were built from direct technical observations, evaluation of records at national, district, and local levels, and structured interviews with scheme users and administrators. Redundant information sources were used to triangulate and assess confidence in our factual findings, which were formalized in novel node-network models of each scheme-village pair. Hypotheses were developed and mapped onto these networks via process tracing. We theorize that resilience of a scheme is dependent on healthy, positive feedback loops within these social-technical networks. We also took early steps toward typological theories on scheme failures.

In the course of the Poshiir study, we developed a method for rigorously and objectively reconciling data from disparate sources. For a number of queries on a given topic, the method takes the array of agreements and disagreements between all sources as inputs; it generates most-probable values for the validity of each source's response, as well as mean validities for the sources themselves. Early tests of the method with a verifiable dataset provide meaningful results, and application to the Poshiir study have returned consistent scores when assessing comparable cases. While further investigation is certainly required, the method appears promising.

We also sought to improve the effective flow rate for capacitive deionization (CDI), a treatment technology for removing ionic species from brackish water. We first proposed two novel powering arrangements that would accomplish this by increasing effective ion mobility: *optimal powering profiles* and *pulse-charged CDI*. We developed a theoretical model and governing equations that would allow us to measure internal ion drift rates from externally observable variables. We then designed, built, and tested a prototype CDI cell capable of operating in either mode. We performed experimental work, as well as numerical simulations for pulse-charged CDI over a range of realistic conditions, capturing a set of crucial timescales. We find that practical restrictions required to avoid redox reactions in pulse-charged CDI ensure that optimal powering profiles will always offer the greater benefit.

Providing safe drinking water access in the developing world will require a new generation of technologies. These must be paired, however, with an improved understanding of how to implement such projects, such that infrastructure proves resilient and impacts prove permanent. I am optimistic that the work presented here may contribute, incrementally, to each of these efforts.

To Them

The inescapable consequence of civilization is that the accomplishments of anyone must rightly be called the accomplishments of everyone. Uncountable people have worked to create the structures, environments, and institutions that make studies such as this possible. I can neither print their names here nor even know them all, but for their invaluable contributions, I dedicate this work to them.

Contents

Contents	ii
List of Figures	iv
List of Tables	xiv
1 Introduction	1
2 Why Do Rural Water Projects Fail?	3
2.1 Introduction	3
2.2 Motivations & Background	4
2.3 Study Design	6
2.4 Data Collection	10
2.5 Data Analysis	12
2.6 Early Results & Preliminary Conclusions	20
3 Source Scoring	25
3.1 Introduction	25
3.2 Motivations, Goals, & Concept	26
3.3 Mathematical Framework	27
3.4 The Impact of Query Types	29
3.5 Early Validations	33
3.6 Strengths, Limitations, and Considerations	36
3.7 Conclusions	40
4 History of CDI	42
4.1 Introduction	42
4.2 Motivations	42
4.3 Capacitive Deionization: Basics and Applicability	44
4.4 History of Capacitive Deionization	46
4.5 The Case for Optimized Power and Pulsed CDI	54
4.6 Introduction to Pulsed CDI	55

5	Pulsed-Charged CDI Theory	62
5.1	Introduction	62
5.2	Conventions	62
5.3	Theoretical Background	63
5.4	Our <i>Charge-Current</i> Model	67
5.5	A Case for Numerical Simulation of Late-Timescale Effects	78
6	Pulse-Charged CDI Prototype Design	80
6.1	Introduction	80
6.2	Design Requirements and Initial Design	80
6.3	Failures, Diagnoses, and Iterative Design	83
6.4	Current Prototype Design and Proof-of-Concept	90
7	Pulse-Charged CDI Experimentation	97
7.1	Introduction	97
7.2	Experimental Methods	98
7.3	Analysis and Data Processing	101
7.4	Results and Discussion	106
8	Numerical Simulation for PC-CDI Timescales	113
8.1	Introduction	113
8.2	Goals & Approach	114
8.3	Planar Electrode Model	115
8.4	Integrating Porous Electrodes	119
8.5	Results and Discussion	123
8.6	Conclusions	131
9	Conclusions and Closing Remarks	135
9.1	Introduction	135
9.2	Conclusions from the Poshir River Watershed Study	135
9.3	Conclusions from the Pulse-Charged Capacitive Deionization Work	137
9.4	Closing Thoughts	139
	Bibliography	141

List of Figures

- 2.1 At left, a map showing India (yellow, pale orange represents regions claimed by India, but currently occupied by other nations), the state of Maharashtra (green), and the three districts intersecting my study area (violet), Raigad, Thane, and Pune. At right, my study area, the Poshiiir River watershed. The dots represent the locations of the water sources for schemes included in the study. The thick lines are subdistrict borders. The thin lines represent drainage divides. The violet region flows into the main Poshiiir River; the yellow region flows into the North fork; and the green region flows into the South fork. Lightness represents elevation. The map was made in QGIS, with elevation overlays from from India's National Remote Sensing Centre [27]. Source locations are based on GPS readings from the field. 8
- 2.2 The baseline scheme-village network used as a template for scheme-specific networks used in formalizing hypotheses. The blue nodes indicate elements integral to the scheme and within the village. Green nodes are integral to the scheme, but outside the village. Red nodes were part of the initial implementation of the scheme, but not the ongoing operations. Finally, village leadership was included as it was relevant for the village politics factor, although not necessarily integral to the scheme. The links represent those relationships expected to exist, though additional links may be added or removed in scheme-specific networks. Karjat MIO and Raigad ZP represents the Minor Irrigation Office and Zilha Parishad respectively, local government entities involved in implementing such schemes. MSEB is the Maharashtra State Electricity Board, the utility providing electrical power in the area. 14
- 2.3 The network representation of the Ainachiwadi scheme. Note that aside from a committee, there are no independent administration or maintenance structures, so those nodes have been grayed out. Also note that the scheme users node has been split into household heads and water gatherers, as the two subgroups were found to have substantively different relationships. Red links represent relationships in tension, along with our confidence scores for that finding. Red-bordered nodes represent disrupted nodes, along with our confidence scores in those findings. Blue links represent findings of explicitly healthy relationships. 15

- 2.4 The network representation of the Poho scheme. As with the Ainachiwadi scheme, operation, maintenance and finance were managed by a single group, here, a family within the village. The independent administration and maintenance nodes have therefore been grayed out. While the technical assets appear to be in a mixed state, the scheme still provides meaningful water access. Again, red represents relationships in tension; blue represents relationships found explicitly good, or nodes explicitly found to be not disrupted. 19
- 2.5 The network representation of the Baliware scheme, following resurrection by the village women. Here, finance, administration, and water gathering are represented by the same node of women in the village (again, inactive nodes have been grayed out). The dashed lines represent two maintenance paths: one directly from the women to technical assets, and one through fitters from outside the village. The scheme, which has proven robust in this form, has simple feedback loops for repair and electricity, with links either found to be healthy or not found to be in tension. 22
- 3.1 Graphical representation of four different comparisons of query responses. In each case, the two axes represent the likelihood of various responses from the two sources being compared. In each case, the light tan fractional area is the the probability of exactly one accurate response, and the dark tan fractional area is the probability of accurate agreement. In the upper left, there are two inaccurate responses of equal likelihood, from each source; $F_a = 1/2$. In the upper right, there are five inaccurate responses, again with equal likelihoods from each source; $F_a = 1/5$. The lower left shows a case with five inaccurate responses with uneven, but correlated distributions, a combination which elevates $F_a > 1/5$. The lower right shows a case with identical but uncorrelated distributions; in such cases, F_a will tend toward the equal-likelihood value (here, $1/5$) represented above and captured in Equation 3.9. 30
- 3.2 A graph showing results for the Lucky Charms example. The ordinate is validity, estimated via Source Scoring; the abscissa is the fraction of responses found to be correct after the analysis. The red line represents the ideal, linear correlation for this distribution, and the data has an $R^2=0.84$ for this line. However, the data can be more closely fit to a curved line: either trinomial (not shown) or erf (in green), each with $R^2 \sim 0.88$. This can be interpreted as a tendency for the method to overstate the deviation from baseline validity. I note that even the abscissa fractions, while perfectly known, are small-sample surveys of their underlying validity. Nonetheless, there is a clear correlation. 35

3.3	Comparison networks for the Lucky Charms example, at left, and the Chahuchiwadi scheme pilot, at right. The nodes represent sources, and each connection represents a query common to the connected sources. The dashed lines represent absent connections. The left network is not ideal because sources s1 and s5 are only weakly connected. That said, it is much more robust and more uniform than the one at right, which has strongly-connected user interviews (UIA, UIB), but leaves the both the national records (NR) and the local records (LR)-observations (Obs) subnetwork weakly-connected.	38
4.1	A typical CDI configuration. Water flows between two high-surface-area electrodes held at a constant voltage by an outside supply. Upon entering, ions in the solution are drawn to the oppositely charged electrode and held in the electrical double layers that form at the surfaces.	44
4.2	Voltage, current, and effluent concentrations for a typical CDI device with a constant flow rate. The abscissa is time. Voltage and Current are shown relative to zero. Effluent concentration is shown relative to influent concentration. Note that during the removal phase, current initially spikes, but then decays. After an initial break-in period, the shaded concentration integrals for the two phases will be equal and opposite.	45
4.3	Helmholtz model (1853) [63], [64]: ions form a single layer at the closest possible distance, entirely screening the electric field. Ions are treated as point charges, with no limit on ion density.	48
4.4	Gouy-Chapman model (1909-1913) [58], [59], [28]: ion concentration decays exponentially from the closest possible distance, screening the field gradually, consistent with Maxwell-Boltzmann statistics.	48
4.5	Gouy-Chapman-Stern model (1924) [129]: integrates an inner "Stern" layer at the "Helmholtz plane" and a diffuse layer that decays into the bulk.	48
4.6	Model including an inner Helmholtz plane of unsolvated ions and outer Helmholtz plane of solvated ions [60], alignment of polar electrolytes near the interface [21], and a diffuse layer with characteristic thickness of Debye length, λ_D	48
4.7	Prominent CDI geometries. Image a. shows traditional <i>flow-past</i> CDI, in which ion drift is orthogonal to fluid flow. Image b. shows <i>flow-through</i> geometry in which the fluid flows through electrodes of alternating polarities. Here, ions move largely in the direction of flow, albeit either faster or slower than the ambient fluid. Image c. shows <i>electrostatic ion pumping</i> , where traditional cells bridge two flow channels. Only the lower-channel exit is opened during the charged-electrode, ion-removal phase; only the upper-channel exit is open during the regeneration phase. Image d. represents arrangements where the electrodes themselves are moved between removal and regeneration phases. Images partially adapted from Porada [114].	52

4.8	Voltage, current, and effluent concentration profiles for traditional constant voltage CDI (or MCDI) at left, and constant current MCDI. Of note is the more constant effluent concentrations of MCDI.	53
4.9	Comparison of equilibration following instantaneous application of constant voltage (left column) and near-instantaneous delivery of finite charge (right column). In each case, the sequence shows the initial state, a moment in transition, and a near-final state. In both cases, the steady-state condition would require an infinite time. Note that while total potential drop is held constant in traditional CDI, the slope of potential between electrode and double layer (dependent only on electronic charge in the electrode) remains constant in the pulse-charged CDI equilibration. See <i>Pulse-Charged CDI</i> section below.	57
4.10	Circuit (left) and pulse profile (right) for Voltage-pulsed CDI. With each pulse of the supply voltage (blue), the cell voltage (red) responds with a timescale equal to RC , where R is the lead resistance(s), and C is the effective capacitance of the cell at those timescales. When the supply voltage is forced down to the DC baseline, the cell voltage discharges with the same timescale. Redox reactions require a minimum voltage to be thermodynamically possible, and a characteristic time to be instigated. This rise and fall must occur fast enough that the CDI cell voltage is above the redox voltage for less than the instigation timescale.	58
4.11	Circuit (left) and pulse profile (right) for Pulse-charged CDI. At each moment, the cell voltage (red) is responding to the higher of the pulser voltage (blue) and the DC voltage (green). The rise of cell voltage is similar to that of voltage-pulsed CDI, with an RC -timescale. The collapse, however, is driven by an equilibration entirely internal to the CDI cell. This faster equilibrium allows a longer pulse without triggering redox. The parallel-diode structure means that the collapse asymptotic to a value below the DC offset, which further decreases the time required to get below redox voltages.	60
5.1	Our theoretical model of the CDI cell. Two parallel, planar electrodes (tan) a finite distance d_z apart, with substantially larger orthogonal dimensions. When the cell is externally raised to a constant voltage, electrons flow to and from the electrodes. In our model, this electronic charge lies in planes at the interior edges of the electrodes. Ions move toward the electrodes, attracted by this charge. In our model, (net) ionic charge is confined to two planes, each one Debye length away from each interface. Relative electric potentials are then defined by five linear regions. To make the simplifications explicit, the potential curve can be compared with those in Figure 4.9.	69

- 5.2 My circuit models of the CDI cell. The first model (image a.) represents the CDI cell as a modified capacitor. The outermost vertical plates represent the CDI electrodes. Charge on the top half of those plates represent electronic charge in the electrode balanced by electronic charge in the opposite electrode, a parallel-plate capacitor with capacitance C'_1 . Charge on the bottom half of the electrodes represents electronic charge in the electrodes balanced by ionic charge in the adjacent double-layer, supercapacitors with capacitances C'_2 and C'_3 . As the ions can move through the bulk, they are connected via resistor with resistance R'_2 . R'_1 represents lead resistances. In the second, equivalent model (image b.), the CDI cell is represented by two capacitors in series with a resistor and capacitor in parallel. The electronic charge in the electrodes is q_e and the ionic charge in the double layer is $-q_i$. $-q_c$ is the shortfall between the two, the electronic charge un-matched by ionic charge. As the inner plates of the outer capacitors and the plates of the inner capacitors exist in the same location, $-q_c$ is a non-physical charge, always canceled by the q_c on the inner capacitor. 70
- 5.3 Interpretation of the charge-current result. At left, the external current, i_e can be seen as charging both the external and internal capacitors, increasing both q_e and q_c . Superimposed on this current, the ionic current, i_i discharges C_1 through resistor R_2 . At right, an arbitrary charging curve. Each incremental bit of the externally provided charge, $dq_e = i_e dt$, initially contributes to q_c (in blue). Each contribution, however, decays with a time constant $\tau = R_2 C_1$ from the moment of arrival. The red represents the charge that has been transferred via ions, q_i . . 75
- 5.4 Spatially local decreases in overall ion concentration near the double-layer. Time, $t = 0.5\tau_D$ (solid), $1\tau_D$ (dot), $2\tau_D$ (dash), $4\tau_D$ (dot-dash), $20\tau_D$ (long dash), where τ_D is *Debye time*. Image is from Bazant [12]. 77
- 6.1 At left, a conceptual diagram of our initial prototype design, based on work by the Farmer group. Water follows a serpentine path upward through the stack. After each pass, it passes through a gap in the electrode plate to the next level. The stack can be scaled to include any number of electrode plate pairs. At upper right, an isometric projection of the actual electrode plate design. The protruding tab is for electrical connection. At lower right, a photograph of the cell being assembled. The metallic material is the electrode plate, and the black material is the carbon aerogel electrode, which has been affixed to the plate with conductive epoxy. Also visible are 20 nylon threaded rods used to compress the stack. 83
- 6.2 At left, a photo showing the original black EPDM electrode-electrode gaskets with large contiguous areas enclosing the nylon threaded rods. At right, a photo showing the thinner electrode-electrode gaskets, here in white EPDM. Other differences presented here include the move from 20 black nylon rods to 12 stainless steel ones and the change from spacers between every level to spacers spanning header plater to header plate. 85

6.3	Photographs from 21-26 Aug 2014. The upper left image shows turbid effluent observed during a run. The right image shows a large fraction of the black electrode that had become detached, as well as the white material visible underneath it. The silver epoxy lines are visible on the detached piece above. The lower left image shows the pitting visible on the aluminum electrode plate after sanding off the remaining epoxy and electrode material. The four horizontal lines where epoxy had been were visible in the pitting pattern.	86
6.4	Photographs showing localized rising in the black electrode (left) and buildup of an initially unknown white substance near the electrode edges (both images). Photos from 21 Oct 2014.	87
6.5	SEM-EDS results for one sample of the white buildup found at the electrode edges. At left, the SEM image. At right, spectroscopy results showing adsorption-dispersion at $K\alpha$ absorption lines for both oxygen and aluminum. There is also a small peak indicating the presence of carbon.	88
6.6	SEM-EDS results for one sample containing a particle of the white material on the nanofoam electrode. At upper left, the SEM image showing the white particle against the black electrode. At upper right, the spectrum shows peaks for carbon, oxygen, and aluminum $K\alpha$ lines. The lower images show the spatial locations of counts for carbon (left, green), oxygen (center, blue), and aluminum (right, red). The carbon locations are consistent with the electrode material in the background, but the oxygen and aluminum locations cluster around the particle.	89
6.7	A conceptual cross-section of the current, stable design. Changes from the initial concept can be seen by comparison with Figure 6.1.	91
6.8	Data from the first successful multi-cycle run with my CDI prototype. The run was conducted in a semi-batch mode, where effluent returned to the source reservoir. The data represents four cycles, each with a removal phase and a regeneration phase. The top image shows the difference between effluent and influent conductivities. Negative values represent removal; positive values represent regeneration. The middle image shows external current, and the lower image shows applied cell voltage. The abscissa for all charts is volume of water exiting the cell. There was a short break in the applied power at $\sim 700\text{mL}$ which seems to have created an uncharacteristic third cycle. The apparent overshoot in the regeneration phase is thought to be an artifact of the semi-batch method.	93
6.9	Image from Gabelich, et al., [47], showing typical data for traditional CDI cycles, included for comparison with Figure 6.8. It shows concentration and voltage for multiple complete cycles. The data was collected in a closed-loop configuration using 1 L, 0.004 M NaCl solutions at 1.0 V, 100 mL/min flow rate, and pH 7. When drawing comparisons, note our use of relative conductivity, $\kappa - \kappa_0$, and their use of fractional concentration c/c_o , as well as our use of volume and their use of time in the abscissa. We do not understand why c/c_o is reported consistently below 1, when we would expect it to rise above 1 during regeneration phases	94
6.10	Design drawing for CDI Electrode Plate, created in SolidWorks.	95

6.11	Design drawing for CDI Header Plate, created in SolidWorks.	96
7.1	Electronic arrangement for observing the initial ionic shielding of the electric field under constant voltage bias (ACP experiments). Green lines indicate the wires carrying the active signal. Orange lines represent (low current) wires used to observe voltages. Black lines are ground. The small box represents a BNC splitter.	99
7.2	Experimental setup for observing the collapse of voltage following a rapid-charging period and isolation from the circuit (RP). Green lines indicate wires for the active signal. Orange lines represent (low current) wires used to observe voltages. Red lines are wires power delivery. Black lines are ground. The small boxes represent BNC splitters.	100
7.3	Fundamental elements of the charging circuit for pulse-charge tests. When $V_{trigger}$ is low, the mosfet is insulative, the capacitor charges from V_{charge} , and V_{out} approaches ground due to the <i>pull down</i> resistor. When $V_{trigger}$ is high, the mosfet becomes conductive, pulling the high-side of the capacitor to ground, and the (negative) voltage pulse discharges through anything attached to V_{out} , in parallel with the 100 Ω resistor. The +15 volt input is required for housekeeping in the mosfet circuit.	101
7.4	Temporally aligned and merged datasets for source voltage, ϕ_{a-e} (blue), and external current, i_e (red), from active-charging phase measurements with 2.5 $\mu\text{S}/\text{cm}$ solution. Each curve represents $\sim 20,000$ data points, each the average of points from three pulses. The transition between the two cadences is visible at 4 μs . We also note the observed voltage rises almost immediately and holds near 1 Volt until ~ 5 ns; this is roughly the time it takes for the electric field to develop in the ~ 6 feet of BNC cable in the CDI-cell circuit. At that point, the circuit appears as a near-short, suppressing voltage until sufficient charge is transferred.	102
7.5	Correcting the Pearson 6027 current monitor's signal decay. The lower, tan line shows data prior to correction – the decay at $\sim \text{ms}$ is due to the Pearson current monitor. The upper red line shows data after correcting for that decay. As the correction operates on an increasingly small value, the uncertainty rises and we see an erroneous upward trend in the corrected curve. Nonetheless, this correction extends usable data by several orders of magnitude. The lighter red line shows the fit to the corrected data through Equation 7.6. The data shown is from active-charging phase measurements, 2.5 $\mu\text{S}/\text{cm}$ solution.	104
7.6	Confirming the charge evolution with independent source-voltage measurement. The thick blue line represents $\phi_{a-e}(t)$ as captured via the oscilloscope. The thin orange line represents $\phi_{a-e}(t)$ as calculated from charge evolutions, with C_1 , C_2 , R_1 , $q_e(0)$ and L_1 (self inductance of the cables) as fitting parameters. Datasets are from active-charging phase measurements, 2.5 $\mu\text{S}/\text{cm}$ solution.	106

- 7.7 Charge evolutions following an increase in applied voltage from 0 to 1 V for four solution concentrations. 2.5 $\mu\text{S}/\text{cm}$ is nearly pure water, where 3440 $\mu\text{S}/\text{cm}$ is brackish. q_e , in red, shows the electronic charge transferred to the cell. q_i , in green, is the total ionic charge adsorbed in the double layer. q_c , in blue represents the difference between the two: the electronic charge in the electrodes unmatched by ionic charge in the double layer. Electric field in the bulk is proportional to q_c . 107
- 7.8 At left, timescales for early shielding, relative to concentration. The timescale was defined as the time at which $q_i = q_c$, or the time at which the field expected based on external addition of charge has reduced by half. The experimental data (in red) can be fit well ($R^2=0.9995$) to a power law (in blue). The fit exponent of -0.97 suggests a hyperbolic relation, consistent with the theoretical *RC time constant* for R_2C_1 as defined in 5.25. At right, the maximum values of q_c , relative to concentration. The experimental data (red) can be crudely fit ($R^2=0.98$) to a power law (blue). 108
- 7.9 Collapse of cell voltage, ϕ_{b-e} , from 1V to 0V after effective isolation from the charging circuit. Data is shown for 37 $\mu\text{S}/\text{cm}$ (left) and 496 $\mu\text{S}/\text{cm}$ (right) concentrations. In each case, red shows experimental data. The green line represents our first fit, based on a single exponential decay of bulk voltage, ϕ_{d-c} . The blue represents our second fit, which incorporates a term likely tied to ionic diffusion into the porous electrode. The strong oscillations at $\sim 10^{-7}$ s are likely due to under-damped conditions due to minimizing R_1 , the large capacitor in our charging circuit, and the self-inductance of our BNC cables. 110
- 8.1 Image showing one half of the planar electrode cell volume, divided into discrete slices. The other half is a mirror image about the center line. Each half consists of three distinct regions. Ions are excluded from the region closest to the electrodes (blue), typically with a width, $B \sim 0.1$ nm. The next region is divided into l slices of uniform thickness, A . Typically, $l \sim 100$ and $A \sim 0.1$ nm. The central region includes m slices. Their widths expand via a geometric series to fill the remaining space. m is typically ~ 150 116
- 8.2 Cross-sections of the porous electrode (tan), illustrating our three regimes. In the *deep electrode* (left), electric fields are fully shielded by electronic charge in the conductive electrode, and ions move through diffusion only in the porous fraction of electrode. In the *transition*, partially attenuated electric fields induce limited electrophoresis in pore centers. In the *electrolyte* (right), ions move via diffusion and electrophoresis within an electric field which is largely uniform in x and y. . 121

- 8.3 Comparing our simulated time-evolution for electric field with those published in the Journal of Physics [91]. The solid black line shows Morrow's double-layer field; the red line shows our equivalent values. The dashed black line shows Morrow's central (bulk) electric field; the blue line shows our equivalent values. The small deviations are likely due to different functional forms for voltage rise and different slice widths. In both cases, the simulated cell has planar electrodes with $d_z=1 \mu\text{m}$, charged to cell voltage of 0.1V for 10 μs , then brought to 0.0V (shorted) for 10 μs . $T = 298 \text{ K}$, $\epsilon_r = 78.5$, and both ions had the (relatively artificial) diffusivity, $D_i = 2.03 \times 10^{-9} \text{ m}^2/\text{s}$. There was no ion exclusion zone ($B=0$), and the width of the first slice (A) was $\sim 0.3 \text{ nm}$ 124
- 8.4 Varying electrode gap, d_z . At left, decay of the central electric field (due to ion shielding) under a constant cell voltage. From left to right, the d_z was 10 μm (red), 30 μm (orange), 100 μm (green), 300 μm (blue), and 1 mm (violet). In all cases, electrolyte concentration was 2 g/L KCl and applied cell voltage was 1 V. At right, data for shielding timescale, τ_s as a function of electrode gap, d_z . Simulation data (blue points) are well-explained ($R^2 = 0.999$) by a linear fit (green line) with time-axis intercept of 90 ns, or approximately the time of maximum central electric field. 125
- 8.5 Varying electrolyte concentration, c_{KCl} . At left, decay of the central electric field (due to ion shielding) under a constant cell voltage. From left to right, c_{KCl} was 50 g/L (red), 5 g/L (orange), 2 g/L (green), 1 g/L (blue), 500 mg/L (violet), and 50 g/L (gray). In all cases, data has been corrected to an equivalent electrode gap of 1 mm, and applied cell voltage was 1 V. At right, data for shielding timescale, τ_s as a function of electrolyte concentration, c_{KCl} . Simulation data (blue points) were fit ($R^2=0.999$) to a power law with exponent -0.929, suggesting an inverse relationship. 126
- 8.6 Numerical simulations of experimental results in the active charging phase. Simulation cell voltage (dark blue) was programmed to match experimental cell voltage (light blue). External current, i_e , was treated as the independent variable. Experimentally measured current (light red) is well replicated by both planar-electrode simulations (green) and porous-electrode simulations (black), until the last decade. At timescales longer than $\sim 10^4$, however, our experimental current is highly dependent on the Pearson correction (see Section 7.3). At left, the the experimental current is interpreted with our best experimentally determined value for the Pearson timescale (1.62 ms). At right, the experimental current has been adjusted to fit the porous simulation by modifying the Pearson timescale to 1.24 ms. This difference in Pearson timescales is far smaller than the discrepancy between our experimentally determined value (1.59-1.65 ms) and that specified by the manufacturer (0.50 ms). Unfortunately, this means that our experimental data is likely insufficient to meaningfully endorse either simulation at late timescales. 127

- 8.7 Active charging phase concentrations (top row) and electric potentials (bottom row), for the whole electrode gap (left column) and near the negative electrode-electrolyte boundary (right column). In both concentration graphs, cation concentrations are solid, and anion concentrations are dotted. Potentials are relative to the potential at the cell's center. In all charts, data is presented at 1 ps (red), 2 μ s (orange), 20 μ s (green), 200 μ s (blue), and 2 ms (violet). Cell voltage asymptote was 100 mV, electrode gap, $d_z=1 \mu$ m, and concentration of 27.5 mM KCl (2050 ppm TDS). At the upper left, we can observe depletion in the bulk, and the local decreases in concentration first presented by Bazant [12] (compare with Figure 5.4), now integrated with ion densities through the electrode. At upper right, we can see the formation (red-blue) and then stabilization (blue-violet) of the electric double layer. At lower left, we see the decaying slope of ϕ , the driving force for electrons in the bulk. At lower right, we see the growing potential drop across the double layer. 129
- 8.8 Relaxation phase simulations of cell voltage collapse. At left, curves show cell voltage following electrical isolation, after charging for 200 ns (orange), 20 μ s (green), 200 μ s (blue), and 2 ms (violet). These colors correspond to electrical isolation starting at the similarly colored times in Figure 8.7. In all cases, the first collapse is an exponential decay and occurs at a charge-independent timescale ~ 2 ns. The second collapse is slower, has a more-nuanced functional form, and a charge-dependent timescale. At right, curves show electric fields at the center of the bulk, $E_c(t)$ (solid blue), and between the electrode and Stern layer $E_{dl}(t) - E_{dl}(final)$ (dotted), for the blue 200 μ s charging case. This shows that the first timescale corresponds to the collapse of voltage gradient in the bulk, while the second corresponds to a decrease in field at the double-layer, consistent with ionic diffusion into the porous electrode. As with Figure 8.7, simulation was for a cell voltage of 100 mV, electrode gap of $d_z=1 \mu$ m, and concentration of 27.5 mM KCl (2050 ppm TDS). 130
- 8.9 Relaxation phase simulations for concentrations (top row) and electric potential (bottom row), for the electrode gap and electrodes (left column) and near the negative electrode-electrolyte boundary (right column). In both concentration graphs, cation concentrations are solid, and anion concentrations are dotted. Potentials are relative to the potential at the cell's center. Data is presented at 1 ps (violet), 20 ns (blue, potential charts only), 20 μ s (green), 200 μ s (orange), and 2 ms (red) after isolation. This data follows 200 μ s of charging at 100 mV (equivalent to the violet line in Figure 8.8). As before, electrode gap, $d_z=1 \mu$ m, and initial concentration of 27.5 mM KCl (2050 ppm TDS). At the upper left, we can observe diffusion into the electrode, as well as increase of ions in the bulk. At upper right, we can see the shrinking double layer. In the lower graphs, we see the bulk voltage collapses completely by 20 ns (violet to blue), and the double-layer voltage shrinks over the later timescales (blue-red). 132

List of Tables

2.1	Single-village piped-water schemes selected for inclusion in the study. Sources are a , as documented in national records[117], b , as observed by researchers, and c , as observed by researchers and reported by villagers in a preliminary group conversation.	9
2.2	Data expected by source and potential factor. The number of checks signifies the relative scope of data collected for each source-factor combination. Other Records includes district and national databases, as applicable.	11
2.3	Candidate hypotheses and hypothesis fragments for the Ainachiwadi example scheme. Hypotheses 1, 2, 4, and 5 are effectively in competition. Fragment 3 could be added to 1 or 2. Fragment 6 could be added to fragment 3 or hypotheses 4 or 5. We have made explicit the total number of elements and the number of those elements without direct motivations in our findings. For each element motivated by a finding, we have included the confidence score for the finding, and calculated upper limits on the confidence we could ever have in that hypothesis or fragment. Between hypotheses 1 and 2, 1 is preferred due to confidence bounds. Between hypotheses 4 and 5, 4 is preferred due to fewer unmotivated elements.	18
3.1	Three versions of a query about Giuseppe Verdi.	32
3.2	Table showing the queries and sources' responses for the Lucky Charms example. Admittedly ad hoc values used for F_a are also shown.	34
3.3	Data for estimated validity (Source Scoring) and the fraction of accurate responses for the Lucky Charms example. Validities with an * were at the upper limit of the validity space explored.	34
3.4	Comparison of estimated source validities for the Chahuchiwadi pilot scheme and the Zugarewadi-Ainachiwadi scheme. Validities with an * were at the limit of the validity space explored.	36
3.5	An invented dataset where no source agrees with any other on any common query, but where the query type varies. Traditional psychometrics and Source Scoring would disagree on the likely validity of source 5.	39

Acknowledgments

This work would not have been possible without the concerted efforts of my research advisor and chair, Ashok Gadgil, and committee members Isha Ray and David Sedlak. I would also like recognize the guidance and support provided by Kara Nelson, Jim Hunt, Bryan McCloskey, Susan Amrose, Heather Buckley, Chinmayee Subban, and my peers in the Gadgil Group, the Blum Center for Developing Economies, and the Berkeley Water Center.

Among current and former members of Lawrence Berkeley National Lab, I would like to thank Will Waldron, Phil Price, Howdy Goudey, and Jonathan Slack. At the Indian Institute of Technology Bombay, special thanks to Milind Sohoni, Pooja Prasad, and Srirang Sohoni. At Santa Clara University, I must acknowledge Betty Young for her unflagging support, and Ed Schaefer for shaking me out of my complacency.

Financially, this work was supported by the National Science Foundation's Graduate Research Fellowship, the State Department's Fulbright Program, the National Science Foundation's IRES grant, and a project grant from the Development Impact Lab. Additional support was thanks to Lawrence Berkeley National Lab, CERC-WET, and the UC Berkeley Civil and Environmental Engineering Department.

I would like to offer a special acknowledgment to my collaborator Jim Keller, without whom much of this could not have happened.

Finally, I must express my profound thanks for the guidance, support, and love received from my family, especially my wife Ashley, who has been infinitely patient with my work, and my young son Jasper, who has convinced me that I shouldn't be.

Chapter 1

Introduction

My Motivations and Approach

When I applied to PhD programs, I sought to gain knowledge, experience, and perspective that I could use to address challenges facing the developing world. To refine my focus, I developed four criteria for potential challenges. They must be problems that are 1) widespread, 2) protracted, 3) where technical interventions might serve to improve conditions, and 4) where those affected were unlikely to resolve the issue on their own. Sadly, many global problems meet all of these criteria.

For my PhD work, I have chosen to focus on drinking water challenges in the rural developing world. Certainly, solving such problems often requires new technology, as limited infrastructure and financial limitations can effectively exclude existing developed-world solutions. New technology, however, while frequently necessary, is often not sufficient to provide sustainable relief. Indeed, even mature technologies frequently fail in these settings. Therefore, even armed with new technical solutions, there is a need to understand the other factors affecting success for such infrastructure projects.

To better understand both aspects, I have chosen to pursue two distinct, but complementary, research paths. The first involves understanding what is required for successful water interventions in rural, developing-world settings. Here, I am developing a set of comparative case studies to explore factors which may affect the success of such programs. The second involves development of improved water-treatment technologies. Here, I am looking to develop a new, more effective type of capacitive deionization that may provide low-cost relief from water stress in regions with brackish groundwater sources.

Guide to Chapters

Chapters 2 and 3 focus on the first research path, presenting the Poshir River Watershed Study, a collection of comparative case studies focusing on single-village, piped drinking

water development projects in rural Maharashtra, India. The study uses a combination of technical field observations, field methods traditionally used in social science, and secondary analysis to build cases for both successful and failed projects, with an aim to developing and prioritizing hypotheses. Specifically, Chapter 2 presents the goals, field methods, and conclusions of the study. Chapter 3 presents a novel analysis method developed for resolving the sometimes contradictory information gathered from disparate sources in the study.

Chapters 4, 5, 6, 7, and 8 concern capacitive deionization (CDI), a treatment method for removing ions from water. Specifically, these chapters focus on the development of a new operational mode we have termed *pulse-charged capacitive deionization (PC-CDI)*. If viable, this would increase effective ion mobility, allowing the same CDI cell to treat more water, thus reducing effective capital costs. Viability, however, depends on understanding key timescales to inform operational parameters, which were not yet understood.

Specifically, Chapter 4 outlines the history of CDI and makes the case for our PC-CDI efforts. In Chapter 5, we present our theoretical work to be used to interpret experimental results. Chapter 6 describes the development of our prototype device and experimental testbed. In Chapter 7, we present results from our experimental determination of early relevant timescales. Finally, in Chapter 8, we describe our efforts to capture latter timescales via numerical simulation, presenting both our model and results showing both confirmation of experimental work and new results.

We conclude the dissertation with Chapter 9. It includes the conclusions of the Poshir River Study, and our assessment of the methods used. It then presents the collective results of CDI efforts, combining theoretical, experimental, and numerical insights, and speaking to the viability of PC-CDI. Finally, I offer my thoughts on the relevance of these paths to the original aim of improving drinking water in the rural developing world.

Chapter 2

Why Do Rural Water Projects Fail?: Building comparative case studies of single-village piped-water schemes in rural Maharashtra

2.1 Introduction

Across the developing world, governments and non-profits are building infrastructure to improve water access for people in rural communities. Such programs have poor success rates, despite recent efforts to identify relevant factors. While such structures face technical failures, the question remains: what factors leave some projects susceptible to those failures while others prove robust? To explore this, we developed in-depth comparative case studies for five working and five failed single-village piped-water schemes in the Poshiiir River watershed of Maharashtra, India. We evaluated potential factors including scheme design, implementation, management, and finances, in addition to source viability, village demographics, and village politics. We accessed this information through a combination of direct observation, examination of records at the national, district, and local levels, and finally through interviews with users and administrators. With only ten schemes in ten villages, this deep-but-narrow study cannot provide firm conclusions. Rather, it will generate hypotheses that can be explored more thoroughly in the future, with shallow-but-broad studies aimed at evaluating specific effects.

While the original goal of this study focused on results surrounding the water schemes being studied, the most important result of this project may be the development of a method for objectively scoring the validity of data from multiple, often conflicting sources, in the absence of a priori answers. When applied to information used in process tracing, this allows us to quantitatively score hypotheses, even within small-n studies. Though employed in

the analysis here, this method, tentatively called *Source Scoring*, is presented in detail in Chapter 3.

In this chapter, we focus on the Poshiir watershed study itself. We first discuss our motivations, describing the challenge and summarizing related efforts found in literature, before formally describing our goals. We then present our study design, addressing our approach and scope, the study area, scheme selection, and dataset creation. We next describe data collection efforts, including study protocols and protection of human subjects. We then describe our data analysis methods, including our representation of the scheme and village as a network, and application of process tracing informed by source scoring. Finally, we present our early results and preliminary conclusions.

2.2 Motivations & Background

While numerous studies have explored factors that might affect the low survival rate of rural drinking water projects, causal mechanisms behind neither failures nor successes are well understood. In this section, I outline the challenge of implementing robust rural water projects in the developing world, briefly review previous efforts to understand such relevant factors, and describe the goals of our study.

Global Challenge

In Maharashtra, India, more than half of communities lack a reliable, sufficient supply of potable water, and more than 1500 communities lack such supplies altogether [25]. Over the last ~50 years, governments and NGOs have undertaken major efforts to provide drinking water to such communities, in Maharashtra and around the world. Increasingly, these efforts have emphasized community participation, decentralization, and investment by the recipient communities (see Kleemeier [75] for a more detailed history). However, success rates for such interventions have proven low, and even now, estimates suggest 30-40% of installed systems worldwide are non-functional [82]. The challenge, therefore, is to understand these failures, such that future projects can be made robust.

Previous Scholarship

Beginning in the 1990s, scholars began to rigorously explore the assumption that community participation led to more successful and sustainable programs (see Narayan [94], Sara [127] broadly, and Manikutty [83] for work in India specifically). While these studies demonstrated a positive correlation, community participation has proven neither strictly necessary nor independently sufficient to create sustainable programs. More recently, scholars have begun to explore a whole constellation of factors, including degree of centraliza-

tion/decentralization [4], gender in management agency [29], transparency [119],¹ degree of locally-informed program customization [10], and sense of community ownership [86]. In addition, the community participation metric has been dissected, focusing on specific elements in planning, decision-making, maintenance, and ownership [119], [25], [82]. For a very thorough history of such participation-focused research, please see The World Bank's recent publication on the topic [85].

The consideration of new factors brought a corresponding increase in methodological breadth. Early studies looking to determine the impact of participation alone often relied on statistical methods for correlation and looked to individual cases to illuminate mechanisms. With the increase in factors of interest, however, comparative case studies have shown promise for identifying relevant factors. Specifically, this can be seen in Manikutty's comparison of two water supply projects in Kerala, India [83], his comparison of five water projects throughout India [84], and Brunner's comparison of 45 projects under two funding programs in Maharashtra [25].

Among these single-factor hypotheses, community participation became fashionable first, and has seen the most scrutiny. These investigations appear to have been motivated by assumption that participation mattered, an assumption which appeared, at least for a time, both ubiquitous and largely unmotivated. The call for decentralization followed a similar track, becoming a popular assumption prior to rigorous research. While subsequent single-factor studies have produced correlations, the initial motivations for evaluating such hypotheses often appear either based in intuition, or reactions to existing assumptions within the field. This has led to two concerns. First, if future scholarship is limited to ungrounded investigations or reactions to dominant assumptions, much effort may be wasted. Second, if methods do not illuminate causal mechanisms, they are unlikely to lead to understanding that can meaningfully inform policy.

Goals of this study

This study seeks to identify potential causal mechanisms behind scheme failures and characteristics that actively contribute to robust schemes, in an attempt to develop candidate hypotheses from a grounded position. To do this, we conducted a deep exploration into a constellation of factors, building comparative case studies of successful and failed schemes. Specifically, we evaluated ten single-village piped-water drinking water schemes implemented in the Poshir River watershed between 2000 and 2011. With the data collected, we can build a collection of comparative case studies with common scheme types, demographics, and hydrological profiles.

¹ Transparency was one of several metrics examined under an umbrella of "good governance," including the more common community participation, and an index termed "project effectiveness."

2.3 Study Design

In this section, we detail our study design. Specifically, we outline early defining decisions, then describe our study area and method for scheme selection. Finally, we discuss the creation of our mixed-source datasets.

Approach & Scope

As the research question focuses on the success or failure of drinking water schemes, we have taken the scheme as our unit of analysis. Formally, our “cases” are then cases of single-village piped-water drinking water schemes implemented by state or district government, drawing water from sources within the Poshir River watershed. While the schemes largely correspond to the village of their implementation, this correspondence is less than perfect.

Several early decisions shaped our approach and the study’s scope. Resource limitations required that we choose between shallow treatment of a large number of schemes and a deeper investigation of a smaller number of schemes. Seemingly comprehensive databases of water schemes are maintained at the national and district levels, suggesting one could approach this research question through a statistical analysis of these data. Alternatively, one can imagine exploring every aspect of a few schemes by going into the villages and observing, measuring, and interviewing. Here, we chose the latter, for three reasons. First, we suspected the relevant elements might include subtle aspects, such as local politics, finances, and relationships not well-captured in high-level databases. Second, we had reservations about the validity of the data in these databases, a concern which has been at least partially vindicated since then. Finally, I was physically in Maharashtra, and the statistical approach could be implemented when I was not.

We also needed to decide whether to select self-similar or highly varying schemes. If we were to select self-similar schemes and find trends,² they would have limited applicability. Conversely, if we cast a wide net, trends would be more applicable, but likely harder to identify. Here we tended toward the former, with the idea that if robust trends could be identified among a relatively self-similar set, later work could explore the applicability of those trends in more varied scheme populations.

Finally, we needed to define the factors to pursue. In contrast to many of the previous approaches, we cast a wide-net here. We included any scheme- or village-specific factors we thought might affect scheme success and which could be captured by the available information sources. Our factors of interest include functionality of the scheme (as a dependent variable), viability of water source(s), state of technical assets, scheme management, scheme

² We are here using “trends” to describe recurring mechanisms or fractions of same. More formally, such patterns might constitute *middle-range typological theories* [53].

finances, scheme maintenance, scheme design, scheme implementation, village demographics, and village politics. The sources identified for each of these candidate factors are shown in Table 2.2.

Study Area

To limit hydrological differences between schemes, we bounded the study within a single watershed.³ The Poshir River watershed (see Figure 2.1) is located ~ 60 km east of Mumbai, largely in the Karjat subdistrict of Raigad district, Maharashtra.⁴ We defined the watershed boundary using a QGIS algorithm and digital elevation map data from India's National Remote Sensing Centre [27].

Geographically, the watershed is bounded by four elevated formations: Peth to the South, Siddhagad to the Northeast, the Bhimashankar hills to the East, and Matheran to the West. The branches of the Poshir flow roughly west until they join the Ulhas River, which flows North and later becomes Vasai Creek, North of Mumbai. The ~ 300 km² region hosts a population of $\sim 40,000$ people [116] in ~ 70 distinct habitations (*wadi*) [120].

In this area, most water is obtained from traditional open wells, small creeks, or the river itself. Deeper borewells are present, but less common, and often built for private use by wealthier farmers. The watershed has spatially consistent subsurface conditions: specifically, a gradually sloping aquitard below 7-12 meters of porous soil. Beginning in late June, monsoons saturate the porous layer and fill the traditional open wells. During the dry season, the river and most open wells run dry before being recharged. As a result, the distances traveled to obtain water increase throughout the season. In recent years, this water stress has been exacerbated by both inconsistent monsoons and growing river pollution.

From 1998 to 2010,⁵ state (*Maharashtra Jeevan Pradhikaran*) and district (*Zilha Parishad*) governments implemented more than 100 single-village schemes (SVSs) aimed at providing drinking water to communities in the region. Most of these SVSs have had no formal impact assessment. Based on early visits, however, we estimated that 60 percent of these schemes had failed. While unfortunate, this presented an opportunity for side-by-side comparison of working and non-working schemes.

³ I here use *watershed* to signify a land area where the surface waters converge. I use *drainage divide* to signify elevated ridges that divide such areas. This convention may be unusual for readers outside of North America.

⁴ The Karjat subdistrict (*taluka, tehsil*) is located in Raigad district (*zilha*). Small parts of the watershed also extend into the districts of Thane and Pune. All are within the state of Maharashtra, India.

⁵ Inclusion in or exclusion from this date range was based on the year in which the scheme was sanctioned, as reported in [117].




Figure 2.1: At left, a map showing India (yellow; pale orange represents regions claimed by India but currently occupied by other nations), the state of Maharashtra (green), and the three districts intersecting my study area (violet), Raigad, Thane, and Pune. At right, my study area, the Poshiiir River watershed. The dots represent the locations of the water sources for schemes included in the study. The thick lines are subdistrict borders. The thin lines represent drainage divides. The violet region flows into the main Poshiiir River; the yellow region flows into the North fork; and the green region flows into the South fork. Lightness represents elevation. The map was made in QGIS, with elevation overlays from from India's National Remote Sensing Centre [27]. Source locations are based on GPS readings from the field.

Scheme Selection

We began scheme selection with a census of all Maharashtra water schemes, taken from India's Ministry of Drinking Water & Sanitation [117]. I then applied the geographic requirement based on our defined study area. By cross-referencing village (*gaw*) locations with watershed boundaries, we identified which schemes potentially drew water from sources within the watershed boundary.

To improve comparability, we then selected only schemes listed as piped-water schemes. We required completion dates before 1 June 2011 to ensure at least one full water year (beginning at monsoon) since scheme completion. We also required that schemes be sanctioned 1998-99 or later, straddling the start of water reforms there.⁶ These criteria reduced the list from 160 to 43 schemes.

⁶ Broadly, the national Swajal Dhara program, and Maharashtra programs Aapale Pani (since 2002) and Jalswarajya (since 2003), emphasize participatory, demand-driven, and transparent planning [25]. The sanction year requirement was intended to capture only schemes made after these programs, but was made based on inaccurate dates for their effective implementation in the field.

These 43 candidate schemes were prioritized based on the likelihood of drawing from a source within the watershed bounds. We visited the villages, confirming the existence of a scheme, the piped-water classification, the GPS coordinates of the source, and the scheme functionality.⁷ Speaking with groups of villagers, we confirmed the rough scheme completion date, implementing agency, and perceived functionality of the scheme. This was continued until we had qualified ten schemes, originally classified as five working and five non-working schemes. A list of the schemes with their initial classifications is given in Table 2.1.

Habitation (<i>wadi</i>) Name ^a	Source ^b	Functional ^c	Sanction Year ^a	Completion Date ^a
Ainachiwadi	open well	Yes	2000-01	12/4/2001
Baliware	open well	Yes	2002-03	3/30/2006
Gorewadi	open well	Yes	2002-03	3/30/2006
Pohi	borewell	Yes	1999-00	11/20/2000
Khairpada	borewell	No	2000-01	3/31/2001
Kalamb	open well	No	2006-07	3/15/2007
Bhaguchiwadi	borewell	Yes	2000-01	3/30/2001
Haryachiwadi	open well	No	2007-08	12/15/2008
Navsuchiwadi	open well	No	2008-09	3/31/2011
Chahuchiwadi	open well	No	1998-99	12/11/2000

Table 2.1: Single-village piped-water schemes selected for inclusion in the study. Sources are *a*, as documented in national records[117], *b*, as observed by researchers, and *c*, as observed by researchers and reported by villagers in a preliminary group conversation.

Dataset Creation

For this effort, we envisioned collecting data via three methods: direct field observation; examination of records at national, district, and local levels; and interviews with users and administrators. Defining the dataset then involved three steps. First, we specified the factors that we sought to capture. Specifically, these included functionality of the scheme, viability of the source, state of the technical assets, scheme management practices, scheme finances, scheme maintenance practices, scheme design, scheme implementation, village demographics, and village politics. Second, we chose specific data that would describe each of those factors. Finally, we compared those data requirements against each of our data sources, including redundancy whenever possible. This method for developing specific tasks was used for both

⁷ While I had anticipated formulating a formal definition for scheme failure, all those we encountered were quite clear cut: either working close to original design or completely defunct. Unfortunately, further investigation would later blur those lines again.

our research goal and our supporting goals.⁸ This breakout of data-collection tasks then informed protocols for field observations and records work, as well as our instruments for structured interviews. The emphases of each data-gathering task are summarized in Table 2.2.

We chose to pursue relevant information from multiple sources because we assumed data from any source would be of imperfect validity. With multiple sources, we imagined some method for triangulation, which could be used to verify and assign confidence to data, where multiple sources were available. That said, at the time of study design and data collection in the field, we had not yet developed the rigorous latent-variable method described in Chapter 3. While this analysis method was developed with the gathered data in mind, the study remains an imperfect match for the analysis method. As a result, there is room for improvement in designing a more unified study, where dataset creation is also optimized for the analysis method.

2.4 Data Collection

Based on the data we hoped to gather for each potential factor, we developed protocols to guide data collection for each class of task. Considerations for the protection of participants were integrated into these protocols. In this section, we discuss the protocols for each class of task, then address human-subject considerations directly.

Protocols

Direct observations focused on the state of technical assets, source viability, and the functionality of the scheme. We photographed and documented the condition of wells, pumps, meters, pipelines, storage tanks, and standposts. Locations of sources, tanks, and standposts were recorded via GPS. We also recorded the capacities and water levels of both storage tanks and open wells.

Records research focused on scheme design, implementation, history, and finances. Demographic information was also available from the national records. National records were obtained from the National Rural Drinking Water Programme [117]. District records were obtained in electronic form, at the district water offices in Raigad. Local records were obtained in hard-copy at the Minor Irrigation Office in Karjat (town).

⁸ The study had several supporting goals, in service of our research goal. 1) As we took the water scheme as our unit of analysis, and intended to combine data from several sources, we needed to ensure that the scheme could be reliably and uniquely identified within each source. 2) We needed to confirm that each scheme fell into the study bounds. 3) For traceability, we maintained a record of how the research was conducted. 4) Finally, we sought to protect those participating from any adverse consequences.

Potential Factor	Sources				
	Direct Observations	Local Scheme Records	Other Records	User Interviews	Administrator Interviews
Functionality of scheme	✓		✓	✓✓✓	
Viability of water source	✓✓✓			✓	✓
State of technical assets	✓✓✓	✓		✓	✓
Scheme management		✓		✓✓	✓✓
Scheme finances		✓✓		✓✓	✓✓✓
Scheme maintenance	✓			✓	✓✓✓
Scheme design	✓	✓✓		✓	✓
Scheme implementation		✓✓	✓	✓	✓✓✓
Village demographics			✓	✓✓✓	✓✓✓
Village politics				✓✓✓	✓✓✓

Table 2.2: Data expected by source and potential factor. The number of checks signifies the relative scope of data collected for each source-factor combination. Other Records includes district and national databases, as applicable.

Within a village, candidate homes for user interviews were selected via careful randomization, based on a census of village buildings built from satellite imagery. Inhabitants were asked to participate if they self-identified as either a primary water gatherer or the head-of-household. Interviews were conducted in Marathi, with at least one interviewer and one recorder, working through translation. Broadly, users were asked about functionality, seasonality, water use, administration, fees, scheme history, and village culture. We conducted between three and six user interviews per scheme, with a total of 44 user interviews over the ten villages.

We identified administrators through the user interviews and then sought them out directly, interviewing one or two current or former administrators per scheme. Broadly, administrators were asked about condition and maintenance of assets, scheme finances, scheme history, and village culture. For all interviews, consents were sought for the interview. In cases where we sought audio recording, a separate consent was sought for that recording.

Protection of Human Subjects

One of our supporting goals focused on the equal distribution of burdens of the study and the protection of participants. This motivated the random selection of users asked to participate in user interviews, as well as efforts to ensure informed consent, and keep knowledge of responses and participation limited to our research team.

Formally, this study and the corresponding protocols were reviewed and approved by University of California's Committee for the Protection of Human Subjects, protocol ID 2012-05-4388. Broadly, this concerned the user and administrator interviews, especially the collection of audio recordings, as vocal recordings are considered personally identifiable information.

2.5 Data Analysis

Data analysis for this study consisted of three steps. First, we looked at questions asked of multiple sources, using the pattern of agreements and disagreements to determine our answers and the degree of confidence in those answers. This was done by a method developed for this study, which we have named *Source Scoring*. This method is presented in detail in Chapter 3. Second, each scheme and village was modeled as a network, based on a template adjusted for each case. Evidence of disruptions or tensions were represented in the nodes and links, respectively. Third, networks were examined for paths through the network that could lead to the observed outcome, rated by the confidence scores for each piece of information. This can be seen as a variant of *Process Tracing*, modified to incorporate the confidence scores, and represented graphically.

In this section, we present the network model for schemes and their villages, which we will use to track and formalize hypotheses. We then briefly introduce process tracing and present our methods for developing hypothesis causal chains, both in failed and in successful schemes.

Network Representation

To transform our descriptive information into a comparable analytical framework, we created a node-network model for the scheme and village. We first created a template network, which was then individualized for each case based on information gathered. These networks were used for tracking causal chains within a case. They also formalize which phenomena will be treated as equivalent in case comparisons. We believe the use of such a network to formalize and visualize hypotheses is novel.

The template network was developed after data collection, but prior to data tabulation and source scoring. It includes nodes based on the factors we sought to investigate, and is shown in Figure 2.2. Broadly, people, organizations, and assets are represented as nodes, and relationships are represented as links. Based on confidence-scored information, evidence for tensions in links and disruptions at nodes, was integrated into the network model. An example of this is shown for the Ainachiwadi scheme in Figure 2.3.

What we have been calling tensions in relationships is a relatively broad category. In the Ainachiwadi scheme example (Figure 2.3), we identified five such unhealthy relationships. The household heads valued the scheme but thought the electricity bills were too high, whereas the water gatherers reported favorable impacts from the scheme. This discrepancy is the basis for our most confident adverse result. The household head view is also the basis for disruption at that node. The two red links originating from the committee stem from our findings that while the committee was well thought of by the villagers (blue link to household heads), they performed no maintenance or repairs on the scheme. The link between the household heads and builders originates from findings of conflict between the groups regarding the use of local labor. Finally, there were reports of unidentified pollutants reaching the source, though that result has a low degree of confidence.

Staying with this example, the disrupted nodes include the potentially polluted source and the technical assets, which suffer from multiple failures, at varying confidence levels.⁹ The design node is considered disrupted due to inconsistent pump specifications in design documents. The household head node is disrupted due to their sentiments about the cost of the scheme's electricity bills.

⁹ While it might be possible to break those assets out into an asset-specific subnetwork, we did not believe we had sufficiently nuanced information about those failures to make it worthwhile.




Figure 2.2: The baseline scheme-village network used as a template for scheme-specific networks used in formalizing hypotheses. The blue nodes indicate elements integral to the scheme and within the village. Green nodes are integral to the scheme, but outside the village. Red nodes were part of the initial implementation of the scheme, but not the ongoing operations. Finally, village leadership was included as it was relevant for the village politics factor, although not necessarily integral to the scheme. The links represent those relationships expected to exist, though additional links may be added or removed in scheme-specific networks. Karjat MIO and Raigad ZP represents the Minor Irrigation Office and Zilha Parishad respectively, local government entities involved in implementing such schemes. MSEB is the Maharashtra State Electricity Board, the utility providing electrical power in the area.

Process Tracing and Hypothesis Generation

Process Tracing Though case studies have been around a long time, formalization of rigorous methods like process tracing and the corresponding epistemological backing is rel-




Figure 2.3: The network representation of the Ainachiwadi scheme. Note that aside from a committee, there are no independent administration or maintenance structures, so those nodes have been grayed out. Also note that the scheme users node has been split into household heads and water gatherers, as the two subgroups were found to have substantively different relationships. Red links represent relationships in tension, along with our confidence scores for that finding. Red-bordered nodes represent disrupted nodes, along with our confidence scores in those findings. Blue links represent findings of explicitly healthy relationships.

actively recent [53].¹⁰ *Process tracing* was coined by Stanford behavioral political scientist Alexander George [52]. It seeks to identify causal mechanisms and causal chains relating independent factors to dependent outcome variables. It does not rely on analysis of statistical correlations, and hence, does not require large numbers of cases. Indeed, it can be used to

¹⁰ The work of John Stewart Mill stands as a notable exception; he published *A System of Logic* in 1843 [90].

determine causal mechanisms within a single case. As such, it is often seen as a complement to large-n statistical methods, which can quantify the scope of correlations, but often do not illuminate causal mechanisms. Process tracing is also frequently used in the development of hypotheses.

In this study, we employ an approach in which *causal process observations* (CPOs) are used to form tests for candidate hypotheses. These tests are of four basic types. *Straw-in-the-wind* tests are neither necessary nor sufficient for the hypothesis but strengthen or weaken it. *Hoop* tests are necessary for a hypothesis, but not sufficient to confirm it. *Smoking-gun* tests are sufficient to confirm a hypothesis but not necessary for it. Finally, *doubly-decisive* tests are both necessary and sufficient for the hypothesis. For more details on this specific approach to process tracing, see Van Evera [37], Bennett [14], or Collier [31].

Process tracing has typically been used in qualitative analysis, but as Collier mentions, CPOs are often based in quantitative result [31]. He warns that care must be taken when classifying such tests; at times, an intermediate test may be in order, neither a hoop test nor a straw-in-the-wind test, for example. Our approach is conceptually similar. We seek to carry confidence scores for CPOs through the application of the test, thus creating a confidence score for the implication(s) of the tests. In this way, hypotheses are not determined to be true or false by smoking gun or hoop tests, but rather the probability of the tested causal chain becomes increasingly bound by the application of such tests.

This is likely not new: as George and Bennett note, “We also find Bayesian logic useful in assessing how ‘tough’ a test a particular case poses to a theory” [53]. Our treatment may be seen as a similar approach, but at a smaller scale, using Bayesian methods to score specific pieces of information and thereby, the tests they pose to a hypothesis within a case.

Hypothesis Generation for Failed Schemes In application, we use process tracing to traverse the networks, generating hypothesis causal chains, which are then scored based on their complexity and their consistency with our findings. In failed schemes, candidate causal chain hypotheses represent pathways, originating with some disrupted node, moving across links, and leading to final disrupted nodes, most typically a failure of technical assets.

Hypotheses can be compared based on the number of steps involved, representing the complexity of the hypothesis.¹¹ Further, the confidence-scored findings of disrupted nodes or links in tension can serve as hoop tests for hypothesis causal chains engaging that element. For example, we could describe a hypothesis for the Ainachiwadi scheme in which heads of households were opposed to the scheme because of high electrical bills, and then refused to pay the committee for its operation. The committee then could not pay the electricity bills

¹¹ A measure of complexity for candidate hypotheses is a crucial metric for abductive inference.

from MSEB, and the utility shut off power to the scheme.¹² The hypothesis is dependent on household heads' frustration with electricity bills. The existence of that disruption then constitutes a hoop test for the hypothesis. Our 80% confidence in that finding then becomes an upper limit for the confidence we might have in that hypothesis. The most feasible hypotheses and hypothesis fragments for the Ainachiwadi scheme are compared in Table 2.5. In this table, our example hypothesis is number 5.

There are two other possible tools for vetting hypothesis causal chains: smoking-gun tests and chronology. In this framework, smoking-gun tests will be less common. Where they do exist, however, they can increase the lower limit of confidence in a hypothesis, again based on the confidence score of the CPO behind the test. Finally, where chronological data exists, this can provide additional vetting for candidate causal chains. Unfortunately, as this data was not collected with process tracing in mind, dates for potential intermediate events were rarely pursued. Where that data exists, however, it can preclude some hypothesis causal chains.

There are several meaningful comparisons that can be made between candidate hypotheses. For hypotheses with equivalent elements, and without unmotivated elements, -we can prefer hypotheses with higher confidence bounds. In the Ainachiwadi example, hypotheses 1 and 2 can be compared in this way. For hypotheses of equivalent confidence bounds, but differing numbers of unmotivated elements,¹³ we can prefer the less complex hypotheses with fewer unmotivated elements. In our example, hypotheses 4 and 5 can be compared in this way. Unfortunately, we have yet to develop a rigorous method for measuring the effect of unmotivated elements on confidence. As a consequence, we cannot present all candidate hypotheses in a single strict hierarchy.

Hypothesis Generation in Successful Schemes For successful schemes, we took a related approach. We acknowledge that a successful scheme could be successful simply because it has, by chance, encountered no disruptions. It could just be lucky. That said, we have assumed schemes successful over many years to be robust, having encountered and surmounted various difficulties. Substantive anecdotal evidence supports this treatment. In analyzing successful schemes, we therefore considered hypotheses for positive response paths to disruptions at the technical asset or water source nodes.

To illustrate this process, we take the Pahi scheme as an example. The network representation of the Pahi scheme is shown in Figure 2.4. Interviews identified past replacement

¹² This example involves many unmotivated steps. We have no evidence that the heads of households refused to pay the committee, that MSEB bills went unpaid, or that MSEB shut off power to the scheme. The example is only meant to illustrate a possible traversal of the network, and the dependence on an observed disruption.

¹³ Such situations will be most common in cases of equifinality where hoop tests only exist for start or end conditions.

#	Hypothesis (H) or Fragment (F)	Total Elements	Unmotivated Elements	Hoop Tests	Confidence Bounds
1	H: Break/theft + Committee did not repair/replace	2	0	92%, 92%	0-85%
2	H: Pollution + Committee did not repair/replace	2	0	55%, 92%	0-51%
3	F: HoH thought scheme was too expensive → HoH told committee they would not pay for repairs	2	1	80%	0-80%
4	H: HoH thought scheme was too expensive → HoH stopped operating scheme	2	1	80%	0-80%
5	H: HoH thought bills were too expensive → HoH told committee they would not pay bills → committee did not pay bills → MSEB shut off electricity	4	3	80%	0-80%
6	F: Inadequate pumping design → Inadequate design led to inadequate scheme with longer pump times → Longer pump times led to higher bills → Higher bills led to HoH thinking bills were too expensive	4	2	47%, 80%	0-38%

Table 2.3: Candidate hypotheses and hypothesis fragments for the Ainachiwadi example scheme. Hypotheses 1, 2, 4, and 5 are effectively in competition. Fragment 3 could be added to 1 or 2. Fragment 6 could be added to fragment 3 or hypotheses 4 or 5. We have made explicit the total number of elements and the number of those elements without direct motivations in our findings. For each element motivated by a finding, we have included the confidence score for the finding, and calculated upper limits on the confidence we could ever have in that hypothesis or fragment. Between hypotheses 1 and 2, 1 is preferred due to confidence bounds. Between hypotheses 4 and 5, 4 is preferred due to fewer unmotivated elements.

of wiring (50% confidence), repairs to two standposts (29%), an electrical transformer replacement (29%), and numerous plumbing repairs (23%) since the inception of the scheme. This scheme appears to be functioning, albeit not as designed, despite likely past failures of technical assets.




Figure 2.4: The network representation of the Pohi scheme. As with the Ainachiwadi scheme, operation, maintenance and finance were managed by a single group, here, a family within the village. The independent administration and maintenance nodes have therefore been grayed out. While the technical assets appear to be in a mixed state, the scheme still provides meaningful water access. Again, red represents relationships in tension; blue represents relationships found explicitly good, or nodes explicitly found to be not disrupted.

At the time of our field work there, the pipe between the borewell and village was leaking underground, losing an estimated 25% of water transferred. In the event of such technical asset failures, there appears to be at least one viable path to repairs, in which technical asset

issues are observed by users, users contact the management family, and the family effects repairs. The feedback loop consists of only three links, two of which were explicitly found to be healthy. The negative finding on the relationship between the assets and users was based on user interview responses, in which some primary water gatherers referred to the leak and the (likely resulting) low pressure. In this case, we do not think it reflects a broad dissatisfaction with the scheme that would lead to disengagement, but rather a desire for improvement indicative of engagement. Indeed, it may be that this sentiment represents the first step in the feedback mechanism we posit.

2.6 Early Results & Preliminary Conclusions

This study is not complete. To date, we have tabulated data for seven of the ten schemes, and performed subsequent full source scoring for four of them. In this section, therefore, we present limited results and preliminary hypotheses. We first discuss our results from successful schemes, then discuss pursuit of trends in candidate hypotheses and hypothesis fractions for failed schemes. Finally, we offer thoughts on the performance of our method.

Characteristics of Successful Schemes

Our analysis methods for both failed and successful schemes rest on the premise that disruptions can traverse the scheme-village network, crossing links through interactions and affecting other nodes, either negatively or positively. For failed schemes, we are comparing hypothesis causal chains connecting early disruptions to the observed negative outcome.¹⁴ For successful schemes, we evaluate possible causal chains connecting presumed disruptions, most often at the technical asset node, to restorative outcomes.

In evaluating successful schemes, we considered the types of issues we had observed in failed schemes. While we had issues with design and construction,¹⁵ we focused on post-construction disruptions, and their possible responses. In effect, we are gauging robustness of functionally designed schemes, or fractions thereof, rather than the viability of initial design and construction. Our hypotheses were then response paths through each scheme-village network.

¹⁴ Disruptions at the technical assets node are perhaps most easily observed, but the Ainachiwadi heads of households' sentiment about about the electrical bills is a similar observation.

¹⁵ In the Baliware scheme, we found that one of six standposts appeared undamaged but proved non-functional. Based on GPS measurements taken at the sites, the standpost is roughly a meter higher than the storage tank apparently meant to supply it. It is likely that this was due to an error in design or initial implementation, and that that standpost had never worked. Similarly, the Ainachiwadi scheme pumping rates and specified operation times were inconsistent with the amount of water meant to be supplied from the scheme. That said, both schemes proved successful, if imperfectly so.

Our tentative theory suggests that a feedback loop capable of effecting repairs, and typified by healthy relationships (links) and invested parties (nodes) in the network, creates robustness with respect to the inevitable failures of technical assets. Robust schemes in Pohn and Baliware both have simple, healthy repair loops. The failed scheme in Ainachiwadi has links under tension and disrupted nodes for any possible response to such events. While we do not believe simplicity is strictly necessary for such loops, our robust cases show loops with few links (and nodes). Loops with fewer elements are less likely to have an unhealthy element. Though less motivated by our data, we posit that a similar loop capable of resolving payment issues may also be necessary.

The scheme in Baliware is an illustrative case. We found that the *gram panchayat*, a local government entity, had previously managed (92% confidence) and maintained (57%) the scheme. It failed in this form, and for two or three years, the women had walked longer distances for water (92%). Then, the scheme was repaired and restarted with the village women in charge of administration (48%), finances (99%), and maintenance (57%). Dissenting sources indicate they contract to “fitters” from outside the village for repair work (57%). It may be that each is true for some fraction of maintenance tasks. The network for the reconstituted Baliware scheme is shown in Figure 2.5. In its current form, which has proven robust for five to six years (86%), it has simple feedback loops for repairs and utility.

In contrast, the original scheme required that an issue with technical assets be brought to the gram panchayat, a path with additional steps. It also likely required involving the household heads, who had mixed responses regarding investment in the scheme.¹⁶ As such, the Baliware scheme stands as a formal example of the successful scheme theory and an informal example of the inverse case.

Should it prove valid after further evaluation, our tentative theory might illuminate underlying mechanisms for correlation-based findings about local management or women’s leadership. The Baliware scheme describes a move from an unsustainable scheme with complex feedback loops under the gram panchayat to a robust scheme when with simple loops under the village women. The Khairpada scheme appears to have a similar trajectory, though we have not yet performed the source scoring for that scheme.

Toward Typological Theories on Scheme Failures

It is our hope that with complete analysis from additional failed schemes, we might compare hypotheses and hypothesis fragments from many schemes. If disruptions move through the network in similar ways, this could constitute a class (type) of detrimental responses

¹⁶While there, I attended a gram sabha meeting where the women gathered in a separate room and watched quietly through an open doorway.




Figure 2.5: The network representation of the Baliwade scheme, following resurrection by the village women. Here, finance, administration, and water gathering are represented by the same node of women in the village (again, inactive nodes have been grayed out). The dashed lines represent two maintenance paths: one directly from the women to technical assets, and one through fitters from outside the village. The scheme, which has proven robust in this form, has simple feedback loops for repair and electricity, with links either found to be healthy or not found to be in tension.

within such networks. That said, our analysis of failed schemes has thus far been less fruitful than hoped.

Chahuchiwadi Of the schemes in this study, the Chahuchiwadi scheme was the first implemented. It was also likely the first to fail. Perhaps as a result, we observe relatively low confidence scores for many of the metrics of interest. Specifically, we know little about the relationships between groups when the scheme functioned. There are also conflicting

accounts of the administration, finance, and maintenance structures of the scheme. As such, constructing a meaningful network is likely futile.¹⁷

Ware The Khairpada scheme began as an effort to bring water from near the habitation of Khairpada both into the habitation and to the much larger village of Ware. This scheme suffered a failure, but was later partially resurrected and modified, delivering water only to Khairpada. Early on, it appeared that the scheme could inform analyses of both failed schemes (through the early Ware-Khairpada failure) and successful schemes (through the resurrected Khairpada scheme). Unfortunately, much of our data on the Ware fraction, including all Ware interviews, has become suspect, and has been deemed unusable.¹⁸ While it may still be possible to obtain meaningful findings from our remaining data on the scheme, we have not prioritized the effort.

Other Failed Schemes There are three other failed schemes that may provide insight. Data from the Kalamb scheme has been tabulated not scored. Data from Haryachiwadi and Navsuchiwadi still require tabulation, most notably interpretation of the onerous local scheme records. As discussed in the previous section, the Ainachiwadi scheme has proven the most fruitful. While we are optimistic that we can meaningfully add to the hypotheses developed there (and presented in the previous section), they must stand alone as our failed-scheme results to date.

Thoughts on Our Method

To the best of our knowledge, there are several novel elements in this study. In the Poshiir effort, source scoring allowed for triangulation on data from disparate sources, estimation of objective confidence scores for findings, and additional comparison of candidate hypotheses within each scheme. Without source scoring, we could still have generated candidate hypotheses via process tracing, but they would have limited comparability, based only on their complexity. Beyond the Poshiir study, the method may prove a useful for a much wider array of applications.

The node-network model also appears novel, especially as a framework for process tracing, and may allow us to trace disruptions in a system to additional impacts, whether beneficial or detrimental. In the Poshiir study, this provided a concrete, abstract framework for formalizing and classifying candidate hypotheses. Casting hypotheses withing this abstract

¹⁷ There was anecdotal description of an industrial plant that had been built, which then tapped the same water source, leading to fights over water access. While our direct observations confirm the existence of a functioning industrial plant and additional draw pipes in the well, this narrative never appeared in formal interviews or local records research.

¹⁸ Suspicious coincidences within notes from otherwise disparate interviews, paired with unquestionable lapses of academic integrity on the part of one interview team, have led us to consider their interview notes unreliable.

framework then allows well-defined comparisons of hypotheses and hypothesis fragments between schemes. This formalization may also have application beyond the study.

That said, the study began with much of the method undeveloped. Indeed, data collection was complete before process tracing was selected and before source scoring was created. While source scoring was developed to work with this data, the study design and data collection could have been optimized for source scoring, process tracing, and the hypotheses we are evaluating.

For source scoring, questions could be included to be asked of multiple sources for the purposes of comparison, even if they did not directly correspond to the factors being explored. This could have created more points of comparison, and more uniform degrees of comparisons between sources (see Section 3.6). For process tracing, more focus could be put on dates of events, as this can be invaluable in assessing candidate causal chains. For the types of hypotheses we are considering, we could have asked about the many more relationships that exist in the scheme-village network. While we did pursue many of these relationships, additional focus would allow a more complete picture of relationships within the scheme and village, and less ambiguity in the assessment of our final conclusions. While potential improvements can be made in subsequent efforts, we are optimistic that continued analysis of the remaining schemes can produce well-grounded hypotheses, worthy of further investigation.

Chapter 3

Source Scoring: A method for objectively scoring data from multiple imperfect, overlapping data sources

3.1 Introduction

In Chapter 2, we describe the Poshiiir River watershed study, in which we built comparative case studies of single-village piped-water schemes in rural Maharashtra. The study design included collecting redundant data whenever possible, as it seemed likely that none of the sources could be assumed to be consistently valid. Indeed, there are many cases where two or more sources returned incompatible answers. The intent was to use data from multiple sources to triangulate and vet the information obtained.

In this chapter, we present that method we developed to do this. *Source Scoring* is a novel but nascent *Bayesian, latent-variable*¹ analysis method for estimating the validity of sources and their responses, given numerous imperfect sources queried about intersecting subsets of a set of questions. here, we present the motivations and concept for the method and the mathematical framework. We discuss the impact of query types, which are central to the method, present preliminary validations, and discuss the method's strengths and limitations. Finally, we present our conclusions for application to the Poshiiir River watershed study, and for the validity of the method more broadly.

¹ The application of latent variables in this method is substantively different than the more common applications in factor analysis or principle component analysis, where the existence of latent variables are inferred by trends within data. In this case, the latent variables are assumed, and the values are inferred via Bayesian inference.

3.2 Motivations, Goals, & Concept

In this section, we describe the motivation for developing Source Scoring, outline the goals of the method, and introduce the underlying concept.

Our Motivating Challenge

We developed Source Scoring for use in our comparative case studies of water schemes in the Poshir River watershed (see Chapter 2). In this study, we gathered data on the scheme’s history from field observations, from interviews with both users and administrators, and from government records at several levels. We included these disparate sources in an intentional effort to build redundancy, and so asked comparable questions of each source whenever possible. The result was datasets where various sources agreed on some points and disagreed on others. While we had intended to use such comparisons to vet our results, we had not identified a method for doing so. Approaching the challenge in a near vacuum, we developed the method described below. Comparisons to adjacent methods in psychology, education, and medicine and were made afterward.

As we developed this method after data collection was complete, the method is designed to be compatible with the Poshir data, but data collection was not optimized for use with this analysis method. Therefore, there are likely ways that future studies could improve on this execution. That said, we are optimistic that the method could be applied to a wide variety of studies.

Goals & Approach

Source Scoring seeks to estimate the validity of data obtained from multiple sources of unknown quality, in cases where multiple sources have been asked comparable questions. It does so by assuming latent variables representing the strength of each source and the difficulty of each question, and treating the whole pattern of agreements and disagreements as manifest variables, probabilistically dependent on these strengths and difficulties. The multi-dimensional latent-variable space can then be searched for the point most likely to produce the pattern of agreements and disagreements observed. This point represents an estimate of the strength of each source and the difficulty of each question.

As the pattern of agreements and disagreements is the input information to this method, it requires a meaningful number of questions asked of multiple sources, such that answers to these questions can be compared. Practical use of the method also requires that some of the sources provide accurate answers.² As we will see, however, it does not require that any sources agree, or that any particular source provide accurate responses to all questions.

² On several occasions, I have had to clarify that this cannot illuminate truths unknown to all of the sources. Rather, it can make information that is collectively held more accessible.

3.3 Mathematical Framework

Since the method's inception, we have begun to describe the underlying theory more formally. In this section, we define our terms and mathematically describe the critical relationships, making assumptions explicit as they arise.

Definitions

We begin by defining a set of sources, \mathbf{S} ,

$$\mathbf{S} \equiv \{s1, s2, \dots\} \quad (3.1)$$

and a set of queries, \mathbf{Q} ,

$$\mathbf{Q} \equiv \{q1, q2, \dots\} \quad (3.2)$$

We then apply queries to sources. We define query set \mathbf{Q}_i as the set of queries asked of source i such that

$$\forall i, \mathbf{Q}_i \subset \mathbf{Q} \quad (3.3)$$

and, symmetrically, define source set \mathbf{S}_a as the set of sources to which query a was applied, such that

$$\forall a, \mathbf{S}_a \subset \mathbf{S} \quad (3.4)$$

We can then define the following:

- $V_{s_i, \mathbf{Q}}$, the *validity of a source i in \mathbf{Q}* , as the probability that a source will return an accurate response to any single query in \mathbf{Q} .
- A *baseline source in \mathbf{Q} , i* , as a source for which $V_{s_i, \mathbf{Q}} = 0.5$.
- $V_{q_a, \mathbf{Q}}$, the *validity of a query, a , in \mathbf{Q}* , as the probability that a baseline source in \mathbf{Q} will return an accurate response.
- A *baseline query in \mathbf{Q} , a* , as a query for which $V_{q_a, \mathbf{Q}} = 0.5$

Formulation

Until now, we have simply defined several *latent variables*. We do not know the values of these validities, but their definitions do not involve any assumptions. In the next step, however, we engage assumptions about how the two validities interact, and posit that the probability of an accurate response when query a in \mathbf{Q} is applied to source i can be estimated as

$$p_{a,i} \cong \text{erf}'[\text{erf}'^{-1}(V_{q_a, \mathbf{Q}}) + \text{erf}'^{-1}(V_{s_i, \mathbf{Q}})] \quad (3.5)$$

where erf' is here defined as

$$erf'(x) \equiv \frac{1 + erf(x)}{2} \quad (3.6)$$

and erf is the *error function* or *Gauss error function*, typically defined via

$$erf(x) \equiv \frac{2}{\sqrt{\pi}} \int_0^x e^{-x^{*2}} dx^* \quad (3.7)$$

or equivalent functions.

The error function and its inverse are not the only option for combining validities. Without a priori understanding of how such terms combine, all we can require is that such a function map two variables, each of range 0-1, to a third variable of range 0-1. That said, this specific formulation has many desirable traits. First, increasing either source or query validity leads to an increased probability of an accurate response under all conditions. Second, a baseline source validity and a baseline query validity combine to give a (baseline) probability of 0.5. Third, equal and opposite deviations from baseline validity cancel each other, again resulting in the baseline probability. This forces the two effects to act on the same scale. Fourth, deviations from baseline are partially additive: two equal and above-baseline validities combine to create a probability higher than either validity alone. Fifth and finally, use of erf is consistent with case-specific deviation from validities following a Gaussian distribution.

The fourth quality excludes simpler functions like arithmetic or geometric means. However, other functions which monotonically map infinite inputs to finite outputs (for example, arctangent), could be made to work similarly. The fifth quality suggests erf as the choice within this class of functions, although additional research is required on which model best maps to real behavior. Until then, the effects of using erf' specifically are expected to be minor, and may manifest in a tendency in predictions toward either baseline or extreme predictions.

We can further define $a_{a,ij}$, the *agreement* between sources i and j on query a as the likelihood that query a will produce the same response from sources i and j . We then note that

$$a_{a,ij} = p_{a,i}p_{a,j} + (1 - p_{a,i})(1 - p_{a,j})F_a \quad (3.8)$$

where F_a represents the probability that two inaccurate responses to query a will agree. This quantity will be treated in Section 3.4.

Within this formalism, known values for the relevant validities would generate the probability that two given sources will agree for a given query. We are interested, however, in the inverse case, where we have no a priori information about these validities (they are latent variables). Instead, we have queried sources and have a known pattern of agreements and

disagreements (manifest variables), but do not know which responses are accurate. In this case, we must identify the combination of validities most likely to result in our observed pattern of agreements and disagreements.³ This is done numerically.

In this way, the approach is Bayesian. Each point in a many-dimensional validity-space represents a hypothesis within continua of hypotheses. Our numerically resolved maxima is the maximum probability of generating our known pattern of agreements and disagreements, given that hypothesis. As such, the approach is subject to the issue of prior probabilities. That is, taking the point of highest contingent probability as a result is only valid to the extent that each hypothesis is equally likely.

Any solution to this optimization will have one unspecified degree of freedom. A point with $erf^{-1}(Vq_{a,\mathbf{Q}})$ increased (decreased) for all a and $erf^{-1}(Vs_{i,\mathbf{Q}})$ decreased (increased) by an equal amount for all i will have an identical probability of generating each response. As a convention, we look to the asymmetry in the validity definitions and constrain this dimension such that the average validity of the queries involved in the optimization is 0.5.

3.4 The Impact of Query Types

One of two principle advantages of Source Scoring lies in the information that can be gained by differentiating between different types of queries. This differentiation lies in the term F_a . In this section, we describe F_a for different conditions, and then, more qualitatively, describe the relationship between validity and agreement for various query types.

F_a in Theory and Practice

As introduced in Equation 3.8, F_a is the probability that two inaccurate responses to query a will agree. In the simplest example, a boolean query a has $F_a = 1$. To discuss other cases of F_a , it is useful to first introduce the idea of the *breadth* of a query, or the number of responses a query might have, and the *distribution*, or how likely each response might be.

We define, A_a , the *absolute breadth* of query a , as the number of valid responses to a query. We also define B_a , the *functional breadth* of query a , as the number of possible responses to a query from sources in \mathbf{S} . Functional breadth excludes valid responses which would not be returned by sources in \mathbf{S} , which would tend to make B_a smaller than A_a . Functional breadth, however, includes invalid responses that would be returned by sources in \mathbf{S} , which tends to make B_a larger than A_a . There is no universally robust way to determine B_a by looking at

³ Given the potentially large numbers of dimensions involved in the validity space, the fraction of space with slightly smaller probabilities may be sufficiently broad that they collectively represent a comparably-likely origin of the observed pattern. In such cases, this likely informs limits on how much can be learned from the method.

a subset of responses, especially when B_a may be larger than the number of sources queried, $|\mathbf{S}_a|$. That said, when the number of unique responses is substantially smaller than $|\mathbf{S}_a|$, we may have reason to estimate a smaller B_a . So estimating B_a remains an art. In our limited practice, we have taken the lower of absolute breadth and estimates based on the breadth of responses received.




Figure 3.1: Graphical representation of four different comparisons of query responses. In each case, the two axes represent the likelihood of various responses from the two sources being compared. In each case, the light tan fractional area is the the probability of exactly one accurate response, and the dark tan fractional area is the probability of accurate agreement. In the upper left, there are two inaccurate responses of equal likelihood, from each source; $F_a = 1/2$. In the upper right, there are five inaccurate responses, again with equal likelihoods from each source; $F_a = 1/5$. The lower left shows a case with five inaccurate responses with uneven, but correlated distributions, a combination which elevates $F_a > 1/5$. The lower right shows a case with identical but uncorrelated distributions; in such cases, F_a will tend toward the equal-likelihood value (here, $1/5$) represented above and captured in Equation 3.9.

In addition to breadth, we should consider the distribution of responses to a query. If there are four possible answers, but less than 1% of sources would give response 1, then response 1 does not functionally contribute to breadth. For queries with finite breadth, we can describe the chance of a given response as the sum of *motives* for that answer, including the motive for the accurate response due to its known accuracy. In the absence of such response-specific motives, this chance becomes a strictly probabilistic chance that the response is randomly selected. We will refer to cases where different responses have equal total motive as *equal distribution* and cases where responses have unequal total motives as *unequal distributions*.

When querying sources of similar type, they are more likely to have similar motives. When combined with uneven distributions, this increases F_a . Sources of sufficiently different type, however, are less likely to have similar motives, which leads to values of F_a which tend toward those for even distributions. Therefore, if we restrict our comparisons to sources of sufficiently different types, we can assume near-even distribution, at which point, F_a becomes

$$F_a \simeq \frac{1}{B_a - 1} \quad (3.9)$$

These cases are presented graphically in Figure 3.1. Given that F_a has some source dependence, treating it as a property of the query a alone must be taken as an assumption. There may be cases, especially where similar classes of sources are compared, where this assumption must be reexamined.

Thankfully, in cases where sources have uneven distributions (which includes most practical cases), F_a is insensitive to small errors in B_a . If we can capture the likely responses, the less-likely responses contribute to F_a according to probability squared, and so have a small impact. This lends some credibility to our practical approach of taking the smaller of absolute breadth and estimates based on the breadth of responses received.

The Relationship Between Agreement and Validity

A strength of this method lies in its simultaneous consideration of each agreement and disagreement over all queries with comparable responses. Nonetheless, it is illuminating to explore the relationship between agreement or disagreement between two sources on a single query and the estimated validities of those sources. The interaction is complex: for queries with a single accurate response, all accurate responses will agree, but whether inaccurate responses agree depends on both the nature of the query and the similarity of the sources. We again use *motive* to describe anything specific to the source that would favor one response, including knowledge about the accurate response, which we will specifically call *accuracy motive*. With these, we can explore several cases. As examples, we will be considering questions about Giuseppe Verdi, born October 10th.

Q1: Was Giuseppe Verdi born on October 10th?	Q2: On what day was Giuseppe Verdi born?	Q3: On what day was Giuseppe Verdi born?
a) yes	a) January 5th	a) January 7th
b) no	b) March 29th	b) March 29th
	c) October 10th	c) October 10th
	d) November 23rd	d) December 25th

Table 3.1: Three versions of a query about Giuseppe Verdi.

In the case of a Boolean queries, where breadth equals two, there is one accurate response and one inaccurate response (see Table 3.1, Q1). All accurate responses will be in agreement. All inaccurate responses, however, will also be in agreement. Here, agreement implies nothing about the validity of either source, but disagreement implies that one source has higher validity, and the other has lower validity. Unfortunately, disagreement on one Boolean query offers no information as to which source is the more reliable one. A set of only Boolean queries, therefore, has substantial drawbacks, as explained below.

Our second case involves a query with a finite number of possible responses and no motives except the accuracy motive. In the absence of the requisite information, no response is more likely than any other. All accurate responses will be in agreement, but inaccurate responses may not be in agreement. The chance of inaccurate agreement will decrease as the breadth of the query increases. In the second Verdi question (Table 3.1, Q2), the chance that any two inaccurate responses agree is likely near one in three. Here, agreement implies both sources have somewhat higher validity: the larger the breadth, the stronger that implication. Disagreement, however, implies less here. Without additional information, we are hard-pressed to discern the difference between disagreement in which one response is accurate and disagreement in which neither response is accurate. I will note that while useful to consider, this case is rather artificial, as additional motives are the norm.

If there are additional motives, it is important to consider whether these motives are common to the two sources being queried. In Table 3.1, Q3, there may be additional motive surrounding d) because it is Christmas in many parts of the world. Depending on context, a respondent might either guess that we included Christmas as an inaccurate option because of its fame, or conversely, guess that we chose to ask about Verdi because he happened to be born on Christmas. In the case where positive motives are common to both sources, the chance of inaccurate agreement will increase, and therefore, agreement does less to imply source validity. Many well-known effects, including courtesy bias and common misconceptions, fall into this category and can effectively masquerade as accurate responses within this analysis method.

In the case where each source has additional motives, but those motives are different, the chance of inaccurate agreement tends toward the case without additional motives. For example, if we asked Q3 of one American and one Ethiopian, the Ethiopian Christmas motives would apply to answer a), instead of d). The odds of agreement on either of these inaccurate responses would be roughly comparable to those of Q2, with unremarkable dates. The more varied the sources, the fewer motives will be common. Therefore, agreement between sources does more to imply higher validity when the sources are of varied type.

In theory, an open-ended question is of infinite breadth. In the case of infinite breadth and no additional motives, there is no inaccurate agreement and agreement always means that the responses are accurate. That said, such queries are almost impossible in real life. Even when questions are open, motives effectively diminish the number of viable responses. When we consider breadth, we must consider not a theoretical breadth, but a functional one, the number of responses that could reasonably result from querying these sources. Further, if we ask open-ended questions, only to classify them into bins, later considered to be in-agreement, then functional breadth cannot be greater than the number of our constructed bins.

3.5 Early Validations

While we were initially uncertain that this method would produce useful results, some early tests have provided cause for confidence. In this section, we present the Lucky Charms example, an artificial trial constructed with verifiable results, and the first two applications to the Poshiir watershed study.

Lucky Charms Example

Early in the development of this method, we conducted a trial which would generate estimated validities for five sources in the absence of a priori data, which would then be compared with source assessments based on perfect data. The queries referred to the Sept 2016 U.S. selection of marshmallows within the Lucky Charms brand cereal. The sources were three students, a random web page result from a search for "Lucky Charms," and an online image of a Lucky Charms box from an unspecified era. Each source was applied to a subset of eight queries. The queries and responses are shown in Table 3.2.

The analysis was performed with relatively ad-hoc estimations of F_a for each of the queries, as we suspected results would be relatively insensitive to small deviations in those variables. The twelve-dimensional space⁴ was numerically searched for the point with the highest prob-

⁴ Each of eight query validities and five source validities constitutes an independent axis for this space, but we are constrained to parts of the space where the average query validity is 0.5, thus reducing it from thirteen to twelve dimensions.

Queries	F_a	Sources				
		s1	s2	s3	s4	s5
q1: Are there yellow charms?	1	-	Y	Y	Y	Y
q2: Are there red charms?	1	-	-	N	N	Y
q3: Are there trees?	1	-	-	-	N	Y
q4: Are there clovers?	1	Y	-	-	-	Y
q5: What color are the diamonds?	1/3	blue	blue	-	-	-
q6: What color are the stars?	1/4	yellow	orange	yellow	-	-
q7: What shape are the purple charms?	1/6	hearts	horseshoes	hats	NA	-
q8: What shape are the white charms?	1/6	circles	NA	circles	NA	NA

Table 3.2: Table showing the queries and sources' responses for the Lucky Charms example. Admittedly ad hoc values used for F_a are also shown.

ability of generating the pattern of agreements and disagreements we had observed, and a (likely) global maximum was identified. The top line of Table 3.3 shows the estimated source and query validities, based on the likely global maximum probability.⁵

Until this point, we had not pursued accurate answers for the queries outside of this method. After generating estimated validities, however, we looked at the Lucky Charms website and treated that data as authoritative. This authoritative data is shown in the bottom line of Table 3.3, and Figure 3.2 plots the validity estimates against the accurate response rate. As these rates are taken to be (small) sample estimates of a larger population, they are subject to standard error with student's t uncertainty.

	s1	s2	s3	s4	s5	q1	q2	q3	q4	q5	q6	q7	q8
estimated validity	0.19	0.89*	0.22	0.70	0.89	0.95*	0.18	0.19	0.95*	0.95*	0.14	0.11	0.45
fraction accurate responses	0.40	1.00	0.20	0.60	0.80	1.00	0.33	0.50	1.00	1.00	0.33	0.25	0.53

Table 3.3: Data for estimated validity (Source Scoring) and the fraction of accurate responses for the Lucky Charms example. Validities with an * were at the upper limit of the validity space explored.

While estimated validities from this method do not show a perfect agreement with the fraction of accurate responses, we do observe a non-negligible correlation. Unfortunately,

⁵ The implicit assumption used in this example is that all points in that space have equal prior probabilities, but this is not necessarily a requirement of the method.




Figure 3.2: A graph showing results for the Lucky Charms example. The ordinate is validity, estimated via Source Scoring; the abscissa is the fraction of responses found to be correct after the analysis. The red line represents the ideal, linear correlation for this distribution, and the data has an $R^2=0.84$ for this line. However, the data can be more closely fit to a curved line: either trinomial (not shown) or erf (in green), each with $R^2 \sim 0.88$. This can be interpreted as a tendency for the method to overstate the deviation from baseline validity. I note that even the abscissa fractions, while perfectly known, are small-sample surveys of their underlying validity. Nonetheless, there is a clear correlation.

the small number of queries per source (and sources per query) limit the resolution of this confirmation. In future validation efforts, we intend to use larger combinations of queries and sources to establish strong baselines, even if the Source Scoring method is applied to only a small fraction of them.

Poshiir Watershed Study Agreements

Early results from the Poshiir watershed study also appear to offer a confirmation of the method. The first two cases analyzed were the Chahuchiwadi scheme (our pilot) and the Zugarewadi-Ainachiwadi scheme. Each of these cases incorporated our field observations, national-level records, district-level records, local (taluka-level) records, and user interviews. While these source types were common, the data for each case was entirely unique. When we found estimated validities for each source, however, the classes of sources scored similarly (see Table 3.5).

Source	Chahuchiwadi Pilot	Zugarewadi-Ainachiwadi
Our field observations	92*%	92*%
District (zilla-level) records	92*%	92*%
User Interview A	92*%	92*%
User Interview B	53%	92*%
National records	72%	82%
Focus Group	na	80%
Local (taluka-level) records	8*%	69%
Administrator Interview	na	55%

Table 3.4: Comparison of estimated source validities for the Chahuchiwadi pilot scheme and the Zugarewadi-Ainachiwadi scheme. Validities with an * were at the limit of the validity space explored.

Despite separate data, the source validities for the two schemes are largely compatible. In both cases, our field observations are estimated at the upper limit of explored validity. District records and user interviews also scored high. National records were a bit worse, and local records were much worse. The two notable discrepancies include the scores on local records and the low score on one user interview, but each discrepancy has an external explanation. The local records are effectively binders of documents about the scheme, and are curated by different employees at the minor irrigation district in Karjat – we do not expect the quality of these to be particularly consistent. As for the user interview, we looked back at the field notes for that interview, and found the entry “Srirang is concerned that he [the respondent] might be drunk.”⁶

3.6 Strengths, Limitations, and Considerations

While Source Scoring appears to hold promise, the method makes several assumptions. Practical application, therefore, requires careful treatment of several factors. In this section, I present the advantages, limitations, and considerations, as we understand them, before comparing Source Scoring to adjacent methods.

Strengths

Source Scoring has several key advantages. First, it provides a method for generating scores for sources, queries, and responses in the absence of a priori information about true answers, and from sources with imperfect knowledge about the subject. Second, it can integrate sources and queries of vastly different types within the same mathematical framework. Third,

⁶ Srirang Sohoni was the Marathi-English interpreter for that interview.

by engaging specific assumptions and integrating query type, it can reveal more information than similar methods.

Known Limitations

We have also identified several limitations. First, if all sources are sufficiently weak (providing inaccurate responses), the method will indicate that, but be unable to identify accurate responses. Indeed, if only one source is at all strong, the method is unlikely to be able to identify the high-validity source. This is little more than stating that the method cannot create accurate information where none existed.

Similarly, if sources are unreliable (providing responses of varying quality), we expect very flat probabilities over the validity space, leading to large uncertainties in resulting estimations. This is because while the latent validities must exist as defined, the central implicit assumption in the method is that they are meaningful. Specifically, for Source Scoring to be useful, the inter-source variability in validity scores must be greater than the query-specific deviation from that validity. In effect, the difference in source validities is the signal, and the query-specific deviations from those validities is noise.⁷

Second, if all queries are Boolean, there will two or more hypotheses (points within the validity-space) with equally likely contingent probabilities. This means multiple (likely quite different) solution sets with regard to query and source validity. One can imagine, for example, that a Boolean query for which all sources returned the same answer could be either very easy or very hard. Further, if the comparison network also has weak points (see below), the method may generate multiples of two solution sets.

Considerations

Several aspects can make this process more or less robust. The optimization process seeks $(|\mathbf{Q}| + |\mathbf{S}| - 1)$ values, based on n response pairs that may agree or disagree. Therefore, this input dataset has $< 2^n$ possibilities,⁸ which places a theoretical limit on the precision of any output. Therefore, the utility of the method increases with the number of comparable responses included in the study design.

The structure of the network of comparisons is also relevant (see Figure 3.3). Weakly-coupled sections of the network can behave relatively independently, and are more likely to lead to multiple contingent probability spaces with local maxima. Long chains attached to

⁷There is an analogous phenomena for query validities.

⁸ Transitivity considerations imply that some fraction of the 2^n space is internally inconsistent and therefore must be excluded. This fraction varies with the number of comparisons in agreement and the types of queries.

the greater network at a single point allow for extreme results for the associated validities. There may be other effects not yet identified. Ideal comparison networks should have uniform and high connectivity throughout.




Figure 3.3: Comparison networks for the Lucky Charms example, at left, and the Chahuchiwadi scheme pilot, at right. The nodes represent sources, and each connection represents a query common to the connected sources. The dashed lines represent absent connections. The left network is not ideal because sources s1 and s5 are only weakly connected. That said, it is much more robust and more uniform than the one at right, which has strongly-connected user interviews (UIA, UIB), but leaves the both the national records (NR) and the local records (LR)-observations (Obs) subnetwork weakly-connected.

As we mentioned in the introduction, this method is Bayesian: in the language of Bayesian probability, each point in the many-dimensional space represents a hypothesis, and our calculation represents the probability of generating the observed outcome, contingent on that hypothesis. As such, the treatment is subject to the issue of prior probabilities. That is, taking the point of highest contingent probability as a result is only valid to the extent that each hypothesis is equally likely.

The process can therefore be made more robust with the inclusion of a source of high, well-constrained validity. There are two ways we can incorporate such sources. First, we could bias the analysis by defining a nonuniform prior probability for the validity of that source. Alternatively, we could assume equally likely hypotheses and expect to see maximal contingent probabilities consistent with high validity for that source. This would provide a

Query	Source 1	Source 2	Source 3	Source 4	Source 5
Multiple Choice (a-e)	a	b	c	d	-
Multiple Choice (a-e)	a	b	c	d	-
Multiple Choice (a-e)	a	b	c	d	-
Multiple Choice (a-e)	a	b	c	d	-
Boolean (T/F)	T	-	-	-	F
Boolean (T/F)	-	T	-	-	F
Boolean (T/F)	-	-	T	-	F
Boolean (T/F)	-	-	-	T	F

Table 3.5: An invented dataset where no source agrees with any other on any common query, but where the query type varies. Traditional psychometrics and Source Scoring would disagree on the likely validity of source 5.

source of independent verification of the application of the method. The latter approach is demonstrated in our treatment of our field observations in the Poshir River project.

Comparison to Similar Methods

Psychometrics Psychometrics is a method developed for measurement of psychological factors. It is often used in education to evaluate testing methods, (equivalent to our sources) assessing many students' understanding (equivalent to our queries). As I understand it,⁹ psychometrics takes information from positive correlations between sources, and takes the degree of positive correlation to be an upper limit on validity. If that correlation results in what we have been calling inaccurate agreement, the validity may be lower.

In contrast, Source Scoring takes additional information from negative correlation, which requires understanding the query type, and is most useful when there is variety of query types in the comparison. For example, see the imagined data in Table 3.5. As there is no agreement, traditional psychometric approach would find no positive correlation between sources on any query and, thereby, infer low validities for all sources. Source Scoring similarly sees the disagreement of sources 1-4 on the first four queries as evidence of low validity for those sources, but sees the disagreements with these low-validity sources on the Boolean queries as cause to improve the validity score for source 5.

Additionally, Source Scoring engages additional assumptions. In that regard, psychometrics will stand on firm epistemological ground for a wider variety of cases. If, however, these

⁹ I have only begun to learn about psychometrics in response to the development of the Source Scoring method. While I believe the information presented here is correct, I am not an expert on the topic and do not want to convey more confidence in my understanding than is due.

more rigorous requirements can be met, and queries can be effectively classified, Source Scoring offers additional information. Finally, because psychometrics returns less information, it usually operates on larger datasets and applies a statistical treatment, whereas Source Scoring can function with fewer inputs, provided the assumptions are well-founded.

Evaluation of Medical Assays We have also become informally aware of a method used in assessing the accuracy of medical assays, used when multiple candidate assays exist, but no gold-standard has been identified. In this method, the assays are equivalent to our sources, and they are assumed to have some (latent) accuracy. The patients are equivalent to our questions, but they are treated as being equally difficult. In this way, it is similar to our method on the assay/source axis, but relies on more traditional statistical treatment on the query/patient axis. Unfortunately, our knowledge of this method is only anecdotal, and further research is required to make a more thorough comparison.

3.7 Conclusions

In this section, we discuss our current view to the viability of Source Scoring as a method, both broadly and in application to the Poshiir watershed study.

On the Viability of Source Scoring

Source Scoring provides a mathematical framework which can integrate responses to queries of disparate types, comparing responses from multiple sources of varying and unknown validity. It allows one to use the pattern of agreements and disagreements to generate probabilistic scores for the strength of each source, the difficulty of each query, and the likely validity of each response.

It generates more information than comparable methods for two reasons. First, it extracts information from disagreements, which requires additional specification of each query type. Second, it engages additional assumptions. Most notably, that meaningful latent validities exist for sources and queries. To be meaningful, the inter-source variability in validity scores must be greater than the query-specific deviation from that variability. These additional constraints make the method less general, but more valuable in applicable cases.

Though nascent, Source Scoring shows promise. Early trials with verifiable data suggest positive correlations with objective measures. Likewise, initial application to our study showed largely consistent results from case to case. That said, the method certainly requires further and more rigorous trials to confirm the behavior and illuminate the bounds of its utility.

In Application to the Poshiir Watershed Study

Source Scoring was not developed until after data collection for the Poshiir watershed study was complete. While the method was developed for work with this data, data collection was not optimized for use with this model. As such, this represents an imperfect application of the method, and there is certainly room for improvement.

Specifically, the source comparison networks have weak points, substantively nonuniform degrees of connectivity, or both. In practice, adjusting to this requires excluding some weakly-connected nodes from the optimization and/or further limiting validity range, as poorly-connected segments of the network tend to lead to extreme values.

The inclusion of our own field observations, however, has allowed us to strengthen the analysis. In our application, we maintain a uniform prior probability, but expect the optimization to return high validities for our field observations as a confirmation. For all cases evaluated thus far, these estimated validities have been high, offering some independent validation of the application.

For much of the failure-mode analysis, source scoring informed process tracing, allowing for *smoking-gun tests* and *hoop tests* where confidence scores assigned to events translate to confidence scores for causal hypotheses. While process tracing without source scoring could still have helped identify candidate hypotheses, source scoring allows us to compare the relative likelihood of candidate hypotheses. That said, the ultimate comparisons still suffers from an imperfect assessment of the complexity of these hypotheses, as is required for abductive reasoning. While I believe useful comparisons can come from the combination of source scoring and process tracing, further development is required here.

The study also sought to identify keys to success. On this effort, we have tentatively identified feedback loops which appear functional in robust schemes and ineffective in failed ones (see Section 2.5). In this analysis, confidence scores for evidence of disruptions in these feedback loops are more directly applicable, and can translate almost directly into metrics for the health of that loop. As a result, these may be the most robust results from the Poshiir watershed study.

Chapter 4

Capacitive Deionization History and The Pulse-Charged CDI Concept

4.1 Introduction

In Chapters 2 and 3 focused on case studies of rural water infrastructure projects, and methods used in that study. At this point, we turn our focus to capacitive deionization (CDI). On this subject, we will cover the history of the technology (this chapter), our theory development (Chapter 5), prototype design (Chapter 6), experimental CDI work (Chapter 7), and numerical simulations (Chapter 8).

In this chapter, we provide motivations for this work, describe the basic operation of CDI, and review the development of CDI technology. We then briefly present the case for optimized power and pulsed CDI, introduce pulsed-CDI arrangements.

4.2 Motivations

In this section, we describe the motivations for our work on CDI. Specifically, we discuss threats to traditional freshwater drinking sources and the challenge of addressing the pending shortfall through the vast, readily-accessible sources of brackish groundwater.

Shortfall and Salinization

In the coming decades, many communities will face acute shortfalls in potable drinking water. Currently, some 2.5 billion people rely solely on groundwater to meet their basic daily needs [118]. Unfortunately, an estimated 20 percent of aquifers around the world are being tapped unsustainably [54]. As expected, this is occurring disproportionately in specific regions. For example, 60 percent of India's groundwater resources are expected to

be depleted by 2030 [35]. While such aquifer depletion will contribute to a host of problems, the shortages in potable drinking water will be among the most acute.

Surface water resources also face unprecedented threats from irregular precipitation and glacial retreat, both due to climate changes. Against this backdrop, seawater intrusion, agricultural practices, and contamination from industry are making many existing freshwater sources increasingly saline. Such salinity is often measured in parts per million (ppm) of *total dissolved solids* (TDS),¹ where 500 ppm TDS represents a level above which, humans typically find the water unpalatable. For this reason, the US EPA limits TDS to 500 ppm as a secondary standard [36]. Historically, the terms *brackish* and *saline* have been notoriously ill-defined; I will here follow Freeze and Cherry's convention, defining fresh water as below 1,000 ppm TDS, brackish water as 1,000-10,000 ppm TDS, and saline water as 10,000-100,000 ppm TDS, and brine as above 100,000 ppm TDS [46].

Freshwater shortages, combined with salinization effects, mean communities are increasingly dependent on brackish and saline sources for drinking water. Already, saline or brackish groundwaters are the sole drinking water source in much of Namibia, and salinization is threatening freshwater supplies in Eastern China, the mountain belt of central Asia, and peninsular India [140]. The people drinking such untreated water suffer.

Like many such challenges, this threat will disproportionately affect the poorest among us. While *reverse osmosis* (RO), *electrodialysis* (ED), and other more mature desalination technologies can effectively remove these salts from solution, rural and developing communities often lack the infrastructure or financial resources required to make these solutions feasible. Indeed, mature technologies frequently fail when implemented in these settings.

The Brackish Water Challenge

As of 2009, an estimated 1.1 billion people lived above brackish or saline aquifers [140]. This water can be a source of suffering, as it is for communities in Namibia. Alternatively, it can become a resource, but only with a viable treatment solution. The *brackish water challenge*, then, is to develop a technology that can affordably produce potable water (TDS < 500 ppm) from the abundant, accessible, and growing brackish aquifers, and thus provide new sources of drinking water for those who cannot afford the current alternatives.

Capacitive deionization is well-suited to address this need, but still faces barriers to scaled implementation in such settings, most notably the capital cost of electrode materials per

¹ The total concentration of salts is also frequently given as (ionic) electrolyte conductivity, κ , measured in $\mu\text{S}/\text{cm}$. To first order, each 1 $\mu\text{S}/\text{cm}$ corresponds to 0.7 ppm TDS, though that factor varies with constituent species.

water treatment rate. Through my work, I seek to improve capacitive deionization to increase production rates, thereby decreasing the effective capital cost of this technology.

4.3 Capacitive Deionization: Basics and Applicability

Capacitive deionization (CDI) is a promising, if not fully mature, water desalination technology. In this section, I describe CDI and its applicability to the brackish water challenge.

Traditional Capacitive Deionization

At its core, CDI is the creation of an electric field between two submerged electrodes, used to draw charged species out of a solution and hold them at the electrode surfaces. Most commonly, CDI is used for desalination of brackish water, though contaminant-specific post-treatment of industrial waste-waters is also common. In a typical CDI configuration (see Figure 4.1), influent water flows between two parallel plate electrodes held at a constant voltage difference, creating the electric field. Dissolved ions in the water are drawn to the electrodes and held by electrostatic attraction to the electronic charge at the electrode surface. The layer(s) of charge at the interface are referred to as the *electrical double layer*.




Figure 4.1: A typical CDI configuration. Water flows between two high-surface-area electrodes held at a constant voltage by an outside supply. Upon entering, ions in the solution are drawn to the oppositely charged electrode and held in the electrical double layers that form at the surfaces.

When the double layers become saturated, flow is usually stopped, voltage is shunted (or reversed), and the ions are released into water that becomes a waste stream. Following this regeneration phase, the process can begin again (see Figure 4.2).




Figure 4.2: Voltage, current, and effluent concentrations for a typical CDI device with a constant flow rate. The abscissa is time. Voltage and Current are shown relative to zero. Effluent concentration is shown relative to influent concentration. Note that during the removal phase, current initially spikes, but then decays. After an initial break-in period, the shaded concentration integrals for the two phases will be equal and opposite.

There are several important design considerations. Electrodes must have high-surface areas, good electrical conductivity, sufficient mechanical strength, and electrochemical stability over a range of pHs. Electrodes are typically made of porous carbon materials with high-surface areas to maximize the number of ions that can be removed before regeneration. BET surface areas from 300-3000 m^2/g are commonly reported, though BET surface area does not map well to absolute capacity.² Porosity, pore size distribution, and other factors appear to have an impact, both on overall capacity and removal speeds.

Operational voltage is dictated by two phenomena. When the electrode is charged, ions of opposite charge, or *counter-ions*, are drawn to the electrodes, but ions of similar charge, *co-ions*, already within the pores are expelled. Initially, these effects cancel each other with

² Methods like BET consistently overestimate the usable surface for ion adsorption in CDI. Gabelich estimated the discrepancy as about a factor of ten [47]. To my knowledge, this remains an only partially explained phenomena.

respect to the net removal of ions from the bulk solution. However, the number of co-ions within the electrode is limited. As increasing charge is transferred to the electrodes, the expulsion of co-ions dies off and the cell begins to remove more ions than it adds. This makes higher charge arrangements – and therefore higher cell voltages – more efficient.

Unfortunately, sufficiently high voltages at the electrode interface will induce oxidation-reduction (*redox*) reactions. Depending on the reaction, this could degrade the electrode, plate material to the surface, create gas cavities within the pores, or further affect local pH. In many cases, this represents an irreversible step. In all cases, it represents charge transfer and unrecoverable energy expenditure. Thus, CDI is usually designed to avoid or limit redox reactions, and cell voltages are typically 1-1.4 V.

Applicability to the Brackish Water Challenge

Worldwide, reverse osmosis (RO) is currently the dominant desalination technology. However, as RO requires that water flow through membrane(s), there is a concentration-independent pressure required to overcome the membrane(s), in addition to the concentration-dependent osmotic pressure. The result is energy requirements that are a comparatively weak function of influent concentrations at lower-brackish regimes. In contrast, energy requirements for CDI are a comparatively strong function of influent concentration across all ranges. Therefore, while CDI is unlikely to compete with RO for seawater desalination, the technology appears likely to be energy-competitive in brackish regimes [105].

In addition, the low pressure requirements, low operating voltages, efficiency even at small-scales, and compatibility with intermittent power sources present further advantages to using CDI in small-scale, remote conditions, like those often found in the developing world. Unfortunately, little of the recent CDI research has focused on this application.³

4.4 History of Capacitive Deionization

In this section, we provide a history of the development of capacitive deionization as a technology. Specifically, we describe the development of electrical double layer theory, pioneering CDI work beginning in the 1960's, revitalizing innovations in the mid-1990's, and recent work on applications, electrode materials, physical configurations, and electrode treatments. We also present the limited work with pulsed power profiles. For additional information, see recent reviews by Slavomir Porada [114], Marc Anderson [2], Yoram Oren [105], and most recently, by the *CDI Working Group* collaboration [130].

³ T. David Waite and Peter Kovalsky's group at the University of New South Wales is a notable exception, and work with the Tata Institute on CDI for specific contaminant removal in rural India.

Development of Electrical Double Layer Theory

When a solid is immersed in an electrolyte, we observe an electrical charge on the solid, due to behavior of functional groups on the solid surface, preexisting static charge, or externally driven charging. In response, ions of opposite charge accumulate around the solid, forming what is called the electrical double layer. CDI relies heavily on the nature and dynamics of electrical double layers. As such, I here outline the development of that theory.

Electrical double layers were first described by Herman von Helmholtz in 1853 [63], [64], who modeled a single layer of uniform ion concentration surrounding the solid (see Figure 4.3). In this way, the double-layers store charge and acted as capacitors, in agreement with experimental observations. This did not, however, explain observations showing that double-layer capacitance varies with voltage and ionic strength of the electrolyte.

This was resolved, independently, by Louis George Gouy [58], [59] and David Leonard Chapman [28]. They describe a diffuse layer where ion concentration decreases with distance from the solid surface, consistent with Maxwell-Boltzmann statistics (see Figure 4.4). It would come to be known that the diffuse layer has a characteristic length at the *Debye distance*, λ_D , typically a few nanometers.

In 1924, Otto Stern combined these models, describing a formation with both a tightly-held inner layer and a diffuse layer [129] (see Figure 4.5). This resolved inconsistencies in highly-charged cases, and has become known as the Gouy-Chapman-Stern, or GCS, theory. In this model, the inner layer is known as a *Stern layer* and exists at the *Helmholtz plane*.

Subsequently, David Grahame incorporated Debye's 1929 description of hydration spheres [34], dividing the Stern layer into solvated ions and those that had shed their waters of hydration, delineating an *inner Helmholtz plane* at a distance on order the ion's atomic radius from an *outer Helmholtz plane* at a distance on order the ion's hydrated radius [60].

Then in 1963, Bockris, Devanathan, and Müller explored the reaction of polar solvent molecules—such as those of water—to charge in the solid [21]. They describe alignment of the solvent molecules at the inner Helmholtz plane, which affects the electrical permittivity (or capacitive dielectric constant) of the solvent. This has become known as the BDM model (see Figure 4.6).

Since then, there have been numerous refinements to double-layer theory, but the BDM model will prove adequate for discussion of pulse-charged capacitive deionization. For further information on more recent developments and current models, see Chapter 6 of Bockris and Reddy [22] or Chapter 7 of Newman [95].




Figure 4.3: Helmholtz model (1853) [63], [64]: ions form a single layer at the closest possible distance, entirely screening the electric field. Ions are treated as point charges, with no limit on ion density.




Figure 4.4: Gouy-Chapman model (1909-1913) [58], [59], [28]: ion concentration decays exponentially from the closest possible distance, screening the field gradually, consistent with Maxwell-Boltzmann statistics.




Figure 4.5: Gouy-Chapman-Stern model (1924) [129]: integrates an inner "Stern" layer at the "Helmholtz plane" and a diffuse layer that decays into the bulk.




Figure 4.6: Model including an inner Helmholtz plane of unsolvated ions and outer Helmholtz plane of solvated ions [60], alignment of polar electrolytes near the interface [21], and a diffuse layer with characteristic thickness of Debye length, λ_D .

It is worth noting that the continuous use of the term *electrical double layer*, even as the model developed, has led to some inconsistencies. In Helmholtz [63], [64], the two layers refer to the layer of charge on the solid and the ionic mono-layer. Newman, however, refers to the two layers as the inner Stern layer and the outer Diffuse layer [95].

Early Experimental Work

The first experimental work on CDI was driven by the US Department of the Interior, in search of an electrochemical method for saline water conversion. This research began at Univeristy of Oklahoma with George W. Murphy, D. D. Caudle, et al. At this point, the technique was called *electrochemical demineralization*, and included processes where redox reactions occur at the electrodes. This early work focused on treatment for graphitic and activated-carbon electrodes such that anions and cations would (electro) chemically adsorb. Seminal works from this era include an exploration of chemically treated graphite electrodes [3] and the first mathematical treatment of their observations [92]. See also [20], [26], [93].

Subsequent work was performed by Allan M Johnson and John Newman at the Marquardt Company in the 1970s, and later by John Newman at U. C. Berkeley and his collaborators. The early work focused on improving the removal process and exploring novel geometric arrangements, generating two patents [71], [70]. They also demonstrated preferential removal of divalent ions and introduced a model treating porous electrodes as a network of resistors and capacitors [67]. See also [68], [69]. Later work by Newman and varied collaborators focused on porous electrode dynamics, *flow-through* arrangements [77], [124], and theory development based on removal of mercury [87].

Concurrently, Yoram Oren, Abraham Soffler, et al. began work in Israel, on electrosorption in a *parametric pumping* arrangement [100], [97], [98], [101], [102]. The group also explored using electrosorption to remove metals [102] and bacteria [103]. In the late 1980s, the group focused on factors affecting carbon electrical double layer capacitance, exploring ion type and concentration [104], charge rate [99], distinguishing slow charging from Faradaic processes [106], pH effects [56], and specific reactions with halides and alkali metals [57].

While this work provided a foundation for future efforts, early attempts at viable CDI devices were hampered by limited options for electrode materials. In Oren's review, he notes that Johnson's group discontinued their work, "mainly because of the instability of the electrodes" [105]. Indeed, excepting Oren's group, there was almost no CDI work through the late 1980's.

The Pekala and Farmer Era

In 1989, however, Richard Pekala of Lawrence Livermore National Lab (LLNL) published a method for making a porous, low-density carbon matrix from resorcinol with formaldehyde,

a material termed *carbon aerogel*, similar to previously developed silica aerogels [108]. In the following years, his group developed this class of materials [109], [112], and began to explore their application as electrodes [88], [111], [73].

In 1995, Joseph Farmer, also at LLNL, demonstrated CDI using carbon aerogels [41], [42], [134]. These materials had the large specific surface area, electrical conductivity, mechanical strength, and electrochemical stability required to make CDI a viable technology. Farmer was the first to call the process *capacitive deionization*.

The Farmer group subsequently demonstrated removal of specific ions including carbonate, sulfate, and phosphate [39]; NaCl and NaNO₃ [43]; and NH₄ClO₄ [38]. They explored applications in wastewater treatment [40] and preferential removal of hexavalent chromium species for groundwater remediation [44]. With this group, Christopher Gabelich later explored preferential removal more broadly and application to natural waters [47].⁴ Finally, Tri Tran, leading the group in the 2000s, used symmetric electrodes and switched polarity each cycle, improving the regeneration phase [135], [136], [137].

During this period, the LLNL groups continued materials research on carbon aerogels and carbon foams. They explored the impact of fabrication conditions [133], [47], material property characterization [110], and potential application in supercapacitors and fuel cells [89] [110].

Applications and New Materials

As capacitive deionization became a viable technology, researchers looked to potential applications. Conway and Ayranci explored treating industrial waste waters, both generally [32] and as applied to specific classes of ions [5], [6], [7]. In, 2005 Wedgemoed and Schrutte evaluated CDI for desalination of brackish and sea waters (although they only treated brackish-level NaCl solutions) [141]. By 2008, however, CDI-based desalination was firmly focused on treatment of brackish waters (for example, see [144])⁵. Treatment of industrial waste waters and desalination of brackish water sources remain the dominant applications.

Much of CDI research since the LLNL work has focused on either developing electrode materials or inventing new operational configurations for CDI. Here, we will cover the advanced materials research briefly, as it is only tangentially related to our work. New materials

⁴Theory describing double-layer formation with competing ions would be presented later by Biesheuvel, et al [18].

⁵In 2008, Yoram Oren published his prominent review [105], which includes a comparison between energy requirements for desalination via CDI and reverse osmosis. He argued, convincingly, that CDI will not be competitive at desalinating seawater-concentration influents, but may prove more energy-efficient at brackish concentrations. This is likely responsible for the focus on brackish desalination post-2008.

evaluated as CDI electrodes include *carbon cloth* [6], [126], [125], *carbon felt* [5], [7], *carbon nanotubes* [33], *nanotube-and-nanofibre-films*, [139], [50], [51], [80], [49], *carbon nanotube-chitosan composite* [148], *carbon sheet* [107], *ordered mesoporous carbon* [151] [81], *graphene* [79], [138], and *carbide-derived carbons* [113]. A return to activated carbon was also explored as a low-cost alternative [150]. In parallel, work continued on carbon aerogels [72]. Recently, this led to *hierarchical carbon aerogel monoliths*, which retain amorphous structure, but include both μm - and nm -scale pores and have proven well-suited to CDI [11], [132]. See also Oren's review, [105], for a more complete set of (pre-2008) foreign-language papers on exotic CDI electrodes.

Novel Configurations

In addition to the typical geometry, often termed *flow-by* CDI, other groups focused on novel operational configurations. Matthew Suss revived the *flow-through* arrangement [131], first described by Johnson and Newman [68], [67]. Implemented with hierarchical carbon aerogel monoliths, influent water flows through the μm pores of electrodes of alternating charge (see Figure 4.7.b.). This leads to faster adsorption, but may require more energy due to increased pressure head, as well as more pre-treatment to avoid fouling.

Bourcier, et al. developed *electrostatic ion pumping* [23], in which water flows through two parallel channels, separated by electrodes of alternating charge (see Figure 4.7.c.). When the electrodes are charged, water exits through the outlet for one channel; when the electrodes are discharged, water exits through the outlet of the other. In this way, flow between the channels alternates direction, flowing toward one channel as ions are removed, and flowing toward the other as ions re-enter solution. Similar to parametric pumping pioneered by Oren [100], [97], [101], [102], this configuration features separate outlets for treated and waste streams, and continuous influent flow.

Finally, researchers developed configurations where saturated electrodes are removed from the treatment stream for regeneration and replaced, allowing for more continuous flow. In one version, electrodes are fabricated as arrays of rigid wires, which alternate between removing ions from treated water and releasing them into waste water [115]. In a second version, *flowable-electrodes* are small porous spheres, charged and introduced to the stream. While in the stream, they adsorb ions and are then mechanically separated and discharged in a waste bath [66], [62]. Despite these newly-developed arrangements, the traditional flow-by version remains the most popular configuration.

Membrane CDI, Constant Current, and Electrode Treatment

Another class of advances involves material treatments applied to the electrodes. While the earliest work had focused on chemical treatments to create functional groups optimized




Figure 4.7: Prominent CDI geometries. Image a. shows traditional *flow-past* CDI, in which ion drift is orthogonal to fluid flow. Image b. shows *flow-through* geometry in which the fluid flows through electrodes of alternating polarities. Here, ions move largely in the direction of flow, albeit either faster or slower than the ambient fluid. Image c. shows *electrostatic ion pumping*, where traditional cells bridge two flow channels. Only the lower-channel exit is opened during the charged-electrode, ion-removal phase; only the upper-channel exit is open during the regeneration phase. Image d. represents arrangements where the electrodes themselves are moved between removal and regeneration phases. Images partially adapted from Porada [114].

for anions and cations [20], [3] and therefore, polarity-specific electrodes, the benefits of alternating polarity each cycle [135], [136], [137] had pushed research toward symmetric electrodes. This focus shifted again with the advent of *membrane capacitive deionization* (MCDI) [1], [76], [80], [24], [74], [16]. In MCDI, electrode surfaces are coated with either anion- or cation-selective membranes. This provides three advantages. First, the membranes prevent expulsion of co-ions into the bulk, forcing all charge equilibration to come from sorption of counter-ions from the bulk. Second, electrode regeneration can be driven more quickly by reversing the applied voltage without risk that released ions will become trapped in the opposite electrode. Third, it allows for constant-current operation.

Until recently, CDI cells had been operated with a constant voltage throughout the removal phase. In 2012, Zhao demonstrated that MCDI cells could be operated with current held constant [149]. This *constant current (cc)* mode, when used with MCDI, provides an effluent of near-constant concentration during each phase (See Figure 4.8). It also allows for direct tuning of effluent concentration through the adjustment of current. This would not be true with (non-membrane) CDI-cc as co-ion repulsion means current does not correlate directly to net ion removal.




Figure 4.8: Voltage, current, and effluent concentration profiles for traditional constant voltage CDI (or MCDI) at left, and constant current MCDI. Of note is the more constant effluent concentrations of MCDI.

Researchers have also begun using chemical treatments to add charged surface groups, thus adjusting the zeta potential of each electrode. By doing so, they can increase or decrease the capacity for electrodes to attract and adsorb ions, even in the absence of electrical charging. This was discovered accidentally, as a detrimental, initial spike in effluent concentration which occurred when electrodes had become oxidized [30]. The effect was reversed through intentional treatment to produce an initial boost in removal, termed *enhanced CDI* [145]. In the extreme version, called *inverted CDI*, the entire adsorption phase operates without power, though a voltage must be applied in the regeneration phase [48]. The theory behind these and other iterations is well-explained here [17].

Aside from the advent of constant current CDI, there has been very little research done on the use of time-varying power profiles, despite theoretical motivations (see Section 5.4).

Indeed, the only published work I am aware of regarding pulsed voltages for CDI was conducted by Sharma, et al. in 2014 [128]. They demonstrated up to 1.3x ion removal with 1, 10, and 100 mV, 100 Hz - 10 kHz square-wave pulses superimposed on an otherwise constant voltage of 1.2 V. Unfortunately, little mechanistic explanation for these improvements was offered, and there has been no continuation of this work to date.

4.5 The Case for Optimized Power and Pulsed CDI

In this section, we present the barriers to large-scale deployment of capacitive deionization in developing communities and make the case for faster effective ion mobility, and thereby for further development of optimal power profiles and pulsed CDI.

Present Limitations

Despite its energy efficiency, traditional CDI can be slow, which presents a challenge for application at scale. The traditional configuration relies on electrically-induced ion flow across the channel, which is initially rapid, but slows as double layers adjust to compensate. As ions accumulate at the electrodes, much of the potential drop occurs over the electrode double layers, and ions in the bulk experience decreasing potential gradients. Flow-through configurations and flowable electrodes reduce the distance necessary for ions to travel prior to adsorption, but increase required pressure. MCDI maintains higher ion drift rates, but still does not represent optimal removal. Sharma's pulsing achieved enhanced removal, but is not well-understood. There is certainly room for improvement in the powering of CDI cells.

The viability of CDI in developing communities will depend on cost. The limitations in ion drift rate effectively create a minimum *hydraulic residence time* for treatment of water in a CDI cell. This means that either water must flow through a short channel very slowly (low pressure requirements but faster saturation leads to frequent regeneration) or more quickly but through a longer channel (higher pressure but less-frequent regeneration). Increasing treatment rate, however, requires increasing electrode size (either parallel to or perpendicular to flow). Unfortunately, even with relatively mundane electrode materials (e.g., carbon aerogels), the electrodes represent the greatest material cost to a CDI cell. Therefore, the slow ion drift rate in traditional CDI functionally limits the treated-water production rate per capital cost.

Our Approach

Broadly, our efforts to make CDI more suitable for use in the rural developing world focus on increasing the rate of water treatment per capital cost of the technology. More specifically, my work is focused on increasing this limit by increasing effective ion drift

rates.⁶ An improvement that increases ion removal can reduce the per output capital cost of CDI, making it more viable for developing communities.

Broadly speaking, ion drift rate responds to electric field, which in turn, results from applied voltage. Increasing ion drift rate through higher voltages is limited by potential redox reactions at the electrode-electrolyte interface. This limitation has two implications. First, the power profile should be such that the total voltage drops across the two electrode-electrolyte interfaces are just below the redox voltages throughout the removal phase. Second, we should be able to superimpose short-duration voltage spikes above this limit, so long as the spike is shorter than the time required to instigate redox reactions. Neither of these conditions are currently achieved via current CDI configurations.

Our goal has been to identify charging profiles which result in improved ion motion and ion removal, such that CDI might be operated with greater throughput per length of electrode material. We have approached this challenge with a combination of theoretical formulation (Chapter 5), laboratory experimentation (Chapter 7), and numerical simulation (Chapter 8).

4.6 Introduction to Pulsed CDI

Researchers at Lawrence Berkeley National Lab and I have identified two potential pulsing arrangements for CDI. They are collectively covered by pending patent [122] and referred to as *pulsed CDI*. In this section, we provide a conceptual framework for pulsed CDI, describe both *voltage-pulsed CDI* and *pulse-charged CDI*, briefly address optimizing power profiles. Finally, we describe the timescale parameters required to implement such arrangements.

Conceptual Framework

It is useful to informally imagine what happens following a theoretical, instantaneous application of voltage across the electrodes, which is then maintained. To achieve this, charge (electrons) would be need to be (instantaneously) transferred to the electrodes, and due to electrostatic attraction, the net charges in the electrodes would reside at the inner edges. At the earliest timescales, an electric field extends from these charges into the electrode. For a typical gap, d_z , of $\sim 1\text{mm}$, this takes $d_z/c \sim 3 \times 10^{-15}\text{s}$. At this point, neither the electrolyte nor ions have responded, the whole bulk is electroneutral, and the potential drop is linear (see Figure 4.9, traditional CDI, Initial). At this point, it behaves like a traditional parallel-plate capacitor, with a vacuum dielectric.

⁶ I will note that this could also be achieved by developing inexpensive, but effective electrode materials. For this complimentary approach, see work by my colleague, Chinmayee Subban.

Everything in the intervening space then reacts in a way that would partially counter this electric field. The (polar) electrolyte responds most rapidly, orienting dipoles in opposition to the field. This is consistent with BDM theory described in Section 4.4. With this transition, the effective permittivity transitions from that of free space, ϵ_0 , to that of the solution, $\epsilon_0\epsilon_r$.⁷ It also increases the capacitance (the amount of charge required to maintain the voltage difference). At this point, the system still behaves like a traditional capacitor, but with the higher dielectric constant.

The ion response is slower, but the ions too respond to cancel the electric field, drifting to the oppositely charged electrodes. After some time (see Figure 4.9, traditional CDI, Intermediate), the ions in the double-layers have partially *screened* the electric field. The ionic charge there causes a second inflection point and the field in the bulk is further decreased. After a long time, the ions have almost entirely screened the bulk field, and almost all potential drop occurs over the double-layers (see Figure 4.9, traditional CDI, Near-Equilibrium).

Although cell voltage does not rise instantaneously, this conceptually describes CDI operation. All of these effects occur at timescales following the more gradual voltage increase.⁸ It is crucial to note that the electric field in the bulk, which drives ion motion there, is strongest at the earliest timescales. It is therefore, the most effective time of the removal phase. If the initial charge transfer can occur at faster timescales than the initial ionic shielding, this will be especially pronounced.

Voltage-pulsed CDI

The first pulsing method we considered involved supplying rapid, higher-voltage pulses during the removal phase, superimposed on the constant DC voltage. This could be accomplished by two supplies in series (as shown in Figure 4.10). During the voltage pulse, cell voltage would rise, at a timescale based on effective capacitance of the cell and the lead resistance(s). When the pulse ends, the cell voltage would be driven back to the DC voltage at the same timescales. The pulse would increase total cell voltage above the voltage required to induce redox reactions, but it would be forced down again before substantive reactions could occur (see Figure 4.10). This was conceived by Robert Kosteki and Ashok Gadgil at Lawrence Berkeley National Lab, and is here called *Voltage-Pulsed CDI*.

There are two potential benefits to voltage-pulsed CDI. First, each pulse would provide a large electric field. If the voltage pulse can rise faster than the ions can respond, the larger

⁷ To the best I can discern, this is an oversimplification. Rigorously stated, ϵ_r is a complex number, with real and imaginary components, and has a frequency dependent response. Our theoretical, instantaneous charge would then exist across all frequencies and exhibit conceptually complex behavior.

⁸ Note that we have not yet discussed the impact of a porous electrode, which effectively lengthens the time to the near-equilibrium state.




Figure 4.9: Comparison of equilibration following instantaneous application of constant voltage (left column) and near-instantaneous delivery of finite charge (right column). In each case, the sequence shows the initial state, a moment in transition, and a near-final state. In both cases, the steady-state condition would require an infinite time. Note that while total potential drop is held constant in traditional CDI, the slope of potential between electrode and double layer (dependent only on electronic charge in the electrode) remains constant in the pulse-charged CDI equilibration. See *Pulse-Charged CDI* section below.




Figure 4.10: Circuit (left) and pulse profile (right) for Voltage-pulsed CDI. With each pulse of the supply voltage (blue), the cell voltage (red) responds with a timescale equal to RC , where R is the lead resistance(s), and C is the effective capacitance of the cell at those timescales. When the supply voltage is forced down to the DC baseline, the cell voltage discharges with the same timescale. Redox reactions require a minimum voltage to be thermodynamically possible, and a characteristic time to be instigated. This rise and fall must occur fast enough that the CDI cell voltage is above the redox voltage for less than the instigation timescale.

electric field will be effectively unscreened by ionic response. As a result, the effective ion mobility would be much greater than suggested by the time-average of the applied voltages. In essence, one would get the most-effective, early-timescale response with each pulse.

The second potential benefit involves ion-ion interference. As concentrations increase, per-ion conductivity decreases. In such strong electrolytes, an ion is surrounded by a cloud of oppositely charged ions, which will be pulled in the opposite direction by an electric field. Sufficiently high electric fields, however, can disrupt these arrangements, at which point per-ion conductivity approaches the *diffuse limit*. This field-induced transition was first observed in 1927 [143], [142],⁹ was theoretically explained in 1957 [96], and is called the *first Wien effect*. In 2003, Li, et al. induced it via pulsed voltages [78], outside of a CDI context. If this can be induced in CDI, it would represent a second, non-linear improvement in ion mobility.

Unfortunately, the current requirements for voltage-pulsed CDI represents a substantive hurdle, especially if trying to induce the Wien effect. The rapid rise of the voltage pulse

⁹ Secondary sources suggest additional seminal work is reported in *Physikalische Zeitschrift*, but copies from that era appear to be rare, and I have not confirmed this.

requires a charging current given by:

$$i \approx \frac{CEd_z}{t} \quad (4.1)$$

where C is the effective capacitance of the CDI cell, E is the magnitude of electric field, d_z is the gap between electrodes, and t is the rise time of the pulse. For the Wien effect, $E \geq 2 \times 10^7$ V/m. Even a small cell has $C \geq 1 \times 10^{-8}$ F and $d_z \sim 1 \times 10^{-3}$ m. Avoiding redox likely means $t \leq 1 \times 10^{-7}$ s. The result is $i \geq 2$ kA,¹⁰ a current so large that it would create huge resistive energy losses in the leads, require expensive equipment to generate, and likely require additional safety considerations. Even without pursuing the Wien effect, the (smaller) resistive losses would be paid both in raising and lowering cell voltage. In contrast, pulse-charged CDI offers a way to avoid discharge losses and improve the net effect of each pulse.

Pulse-charged CDI

In *pulse-charged CDI*, (PC-CDI), the rise in cell voltage is driven by the supplied voltage pulse as in Voltage-pulsed CDI, but the fall in cell voltage is instigated by the transition from parallel-plate capacitor to supercapacitor, described below. Practically, this can be accomplished by putting a (rapidly responding) diode in series with the pulsing supply. A DC baseline voltage can then be achieved by adding a DC supply, with its own diode, in parallel with the pulser (see Figure 4.11). As with voltage-pulse CDI, the rise in cell voltage has a timescale equal to RC . The fall in cell voltage, however, occurs much more quickly.

In PC-CDI, the voltage drop is due to an equilibration completely internal to the CDI cell. When the supply voltages are below the cell voltage, the diode blocks charge flow, effectively isolating the cell from the circuit. During this time, the electronic charge on the electrodes remains constant. As ions move, however, the cell then evolves from a parallel-plate capacitor—where the electronic charge on an electrode is balanced by electronic charge on the other electrode—toward a supercapacitor—where electronic charge on each electrode is balanced by ionic charge in their respective double-layer. This evolution comes with a dramatic increase in capacitance. *Capcitance*, C , is defined as

$$C \equiv \frac{q}{V} \quad (4.2)$$

where q is the magnitude of equal and opposite charges held in the capacitor and V is the difference in electric potential across the capacitor. Geometrically, parallel-plate capacitance is directly proportional to the area of the plates, A , and inversely proportional to the distance between the electrodes, d , via

$$C = \epsilon_0 \epsilon_r \frac{A}{d} \quad (4.3)$$

¹⁰ If sustained, a 2 kA current will melt a copper rod the diameter of a pencil.




Figure 4.11: Circuit (left) and pulse profile (right) for Pulse-charged CDI. At each moment, the cell voltage (red) is responding to the higher of the pulser voltage (blue) and the DC voltage (green). The rise of cell voltage is similar to that of voltage-pulsed CDI, with an RC-timyscale. The collapse, however, is driven by an equilibration entirely internal to the CDI cell. This faster equilibrium allows a longer pulse without triggering redox. The parallel-diode structure means that the collapse asymptotic to a value below the DC offset, which further decreases the time required to get below redox voltages.

where ϵ_0 is the permittivity of free space and ϵ_r is the relative permittivity of the medium. During the transition from capacitor to supercapacitor, effective values for both A and d evolve. The effective area increases from the projected, area of the bulk electrode to the tortuous total surface area of the nano-porous electrode, increasing by a factor of $\sim 10^3$. Simultaneously, as electronic charge in the electrode transitions from being balanced by electronic charge in the opposite electrode to being balanced by ions in the double layer, the effective distance, d , decreases from the electrode gap, d_z , to \sim the Debye length, λ_D , decreasing by a factor of $10^5 - 10^6$. This represents an increase in capacitance by a factor of 10^8 or more. For a constant charge in the electrode, the voltage drop will collapse by a similar factor. While the timescale for this voltage collapse was initially unknown, experiments (see Chapter 7) and simulations (see Chapter 8) suggest it is very fast.

There are two principle advantages to pulse-charged CDI, compared with voltage-pulsed CDI. First, the faster fall of the cell voltage allows longer pulses, and, correspondingly, increased impact on ion motion. The faster fall is due to both a lower asymptote made possible by the parallel arrangement of the voltage sources and the faster timescale of the internal voltage collapse. Second, all current experienced by the leads serves to charge the electrodes, and thus contribute to eventual ion motion. In contrast, a voltage-pulsed arrangement transfers many times the charge delivered per-pulse, both into and out of the

cell. Thus, resistive losses in pulse-charged CDI are much smaller than in voltage-pulsed CDI.

Relying on internal equilibration for the falling edge of voltage pulses, however, likely limits the maximum pulse voltage. Timescales for equilibration at constant voltage are nonlinear with respect to applied voltage [91], [61]. If internal equilibration timescales under constant charge are similarly nonlinear, there will be a pulse voltage such that the voltage cannot collapse fast enough to avoid redox reactions. As we will see, this will have to be explored through experimentation or simulation.

Optimal Power Profiles

With or without pulsing profiles, current CDI arrangements do not maximize redox-free ion motion. Traditional CDI operates at a constant cell voltage, low enough to avoid or minimize redox (typically 1.0-1.4 V). cc-MCDI operates with a cell voltage that rises to a redox-based cutoff. Any pair of redox half-reactions is driven not by cell voltage, however, but by the sum of voltages across the double-layers. For periods in the removal phase where non-trivial voltage drops occur over the bulk, cell voltage could be safely elevated without risk of redox. Such a configuration would require understanding the time-evolution of the voltage across the bulk.

While this would increase mean ion mobility in traditional CDI, the removal improvements would occur largely at the beginning of the removal phase, exacerbating the already inconsistent effluent concentrations (see Figure 4.8). For applications requiring consistent effluent concentrations, cc-MCDI may still prove preferable. While it is possible that this optimized CDI effluent might be "smoothed" by tuning a variable flow rate, adding a post-treatment CMFR, or merging effluent from multiple out-of-sync cells, those considerations are beyond the scope of this work.

Requisite Parameters (for PC-CDI and optimal power)

At this point, it is clear that understanding the relevant timescales is of crucial importance to pulse-charged CDI. First, the arrangement can only operate in regimes where the time for pulse rise (via RC-charging) and pulse fall (via internal equilibration) is less than the time to induce redox reactions. Second, if pulses induce intermediate energized states in the lead-up to redox, the time between pulses will need to be sufficiently long that these states relax before the next pulse. Third, the effectiveness of pulses will vary with the timescale for decay of the bulk electric field under constant voltage, relative to any potential pulse-width. Even the optimized power profiles would require understanding the timescales for decay of the electric field. Much of the following work in theory (Chapter 5), experimentation (Chapter 7), and simulation (Chapter 8) is aimed at illuminating these timescales.

Chapter 5

Pulsed-Charged CDI Theory

5.1 Introduction

In Chapter 4, we reviewed the history of CDI theory and presented arrangements for increasing effective ion mobility through non-traditional powering methods. Specifically, we noted the importance of understanding timescales for ionic shielding and collapse of the bulk electric field under constant-voltage, and voltage collapse under constant charge. In this chapter, we present theory we will use to interpret experimental data in Chapter 7.

In this chapter, we first establish conventions, defining terms we will use in cases where field-specific conventions diverge. We then briefly describe existing models relevant to capacitive deionization, as well as our own theoretical models and their applicability to PC-CDI. We close by exploring the requirements for a more thorough analytical treatment and make the case for addressing this challenge through numerical simulation.

5.2 Conventions

Electric potential and field

Much of the following theory will involve *electric fields* and *electric potentials*. While the use of the term "electric field" seems consistent, physicists and electrochemists appear to have different conventions regarding "electric potential." In this work, we follow physics conventions. Specifically, and except where otherwise made explicit, I treat electric potential as the potential energy of a unit charge at a given location relative to that unit charge an infinite distance from everything. Consistently, it is the amount of work required to bring that unit charge from infinity to that point (without acceleration). Embedded in this, there appear to be two considerations. First, the functional energy of an electron in a solid is dictated not only by electrostatics of net charges, but also by the surrounding (electrically net neutral) chemical environment. I use *electric potential*, here denoted ϕ_a , to consider only

electrostatics (electrochemists' *Galvani potential*) and *electrochemical potential*, $\bar{\mu}_a$, when considering environmental effects (electrochemists' *Fermi level*). Second, relative electric potentials are often more useful than absolute values. The definition of the zero-value datum is, therefore, largely convention.¹ I use *voltage*, here denoted ϕ_{a-b} , to denote the difference between any two electric potentials ϕ_a and ϕ_b .

Electric field and potential are related by the following equation:

$$\vec{E} = -\vec{\nabla}\phi + \frac{\partial \vec{A}}{\partial t} \quad (5.1)$$

where \vec{E} is the electric field vector, ϕ is electric potential, \vec{A} is the *electric vector potential*, and t is time. Without time-dependent magnetic fields, however, the second term disappears. As this is a reasonable assumption for most of our cases, I will neglect it, leaving the much simpler form

$$\vec{E} = -\vec{\nabla}\phi \quad (5.2)$$

As such, using either term is sufficient to determine the electric field, but only ϕ uniquely specifies the electric potential. Therefore, I typically write equations with respect to ϕ , and also present them with respect to \vec{E} when it seems more conceptually intuitive.

5.3 Theoretical Background

In this section, we describe previously developed theory relevant to CDI generally, and the pursuit of timescales specifically. We then cover thermodynamic limits for desalination and CDI, work done with the Poisson-Nernst-Planck model, and attempts to integrate the effects of porous electrodes. Finally, we briefly discuss the Donnan equilibrium models.

Thermodynamic Limits

Any desalination effort is subject to thermodynamic limits. As desalination constitutes movement from a relatively disordered state to a relatively ordered state, the change in entropy represents a minimum energy expended. For seawater desalination ($\sim 35,000$ ppm TDS, or ~ 50 mS/cm), this minimum energy is ~ 1.1 kWhr/m³. For brackish water of 4,000 ppm TDS (~ 2850 μ S/cm), this limit is ~ 0.12 kWhr/m³ [2]. In both cases, precise values depend on effluent concentration and *recovery rate*, the ratio of effluent to influent flow rates.

¹ Electrochemists tend to use less-abstract conventions for the reference points on such potentials. I believe that in the conventions of electrochemistry, my choice of datum is called *absolute electric potential*, though that applies to an electrochemical, rather than electric potential. I might offer that my convention would correspond to *absolute Galvani potential*, but I have not seen that term defined.

Biesheuvel applied this to an ideal CDI cycle, including both the removal and regeneration phases [15], and generated a CDI-specific equation for minimum energy required per effluent volume,

$$\frac{W}{V_d} = 2\frac{c_0}{\alpha}\ln\left(\frac{c_0 - \alpha c_d}{c_0(1 - \alpha)}\right) - 2c_d\ln\left(\frac{c_0 - \alpha c_d}{c_d(1 - \alpha)}\right) \quad (5.3)$$

$$\alpha \equiv \frac{V_d}{V_0}$$

where W is reversible work given in multiples of $k_B T$, V_d is (low-salinity) effluent volume, V_0 is influent volume, c_d is effluent concentration, and c_0 is influent concentration, both in units of numbers per volume. Through conservation laws, the (high-concentration) waste stream terms have been eliminated. As defined above, α is recovery rate. As an example, a CDI system taking water from 4,000 to 400 ppm TDS with $\alpha=0.5$ requires ≥ 0.15 kWhr/m³. In contrast, bringing water from 1,000 to 400 ppm TDS with $\alpha=0.5$ would require ≥ 0.015 kWhr/m³. Because of dependencies on concentrations, recovery rates, and temperature, care must be taken in comparing reported energy requirements.

Poisson-Nernst-Planck Models

The behavior of ions in the bulk solution can be well-described by the simultaneous solution of two equations. *Poisson's equation* describes the effect of net charge on electric potential at given location:

$$-\vec{\nabla}^2 \phi = -\vec{\nabla} \cdot (\vec{\nabla} \phi) = \frac{\rho_q}{\epsilon_0 \epsilon_r} \quad (5.4)$$

where ρ_q is net charge density, ϵ_0 is the *permittivity of free space* or the *electric constant*, and ϵ_r is the *relative permittivity*. The *Nernst-Planck equation* describes the motion of ions in the presence of an electric field:

$$\frac{\partial c_i}{\partial t} = \vec{\nabla} \cdot \left[D_i \vec{\nabla} c_i - c_i \vec{u} + D_i c_i \frac{z_i e (\vec{\nabla} \phi)}{k_B T} \right] \quad (5.5)$$

For an ion species i , c_i is the concentration, D_i is the *diffusivity*, and z_i is the charge of i ; \vec{u} is the fluid velocity, e is the *elementary charge*, k_B is the *Boltzman constant*, and T is temperature. The term in brackets is total ion flux, and is due to three terms: diffusion, fluid flow, and *electrophoresis* (movement of ions due to an electric field), respectively. Where the first term represents the spatial change in concentration, the last term represents the spatial change in electrical energy, relative to thermal energy. In many of our cases, the fluid flow term can be neglected as fluid is either static or moving perpendicular to the ion flows of interest.

Simultaneous solution of Poisson's equation (5.4) and the Nernst-Planck equation (5.5) over a space gives the time-evolution of ion concentrations. *Boltzman statistics* or *Gibbs statistics* give that the concentration of ions of species i in a given (here, electrical) energy state is given by

$$c_i(\phi) \propto c_i e^{-\left(\frac{z_i e \phi}{k_B T}\right)} \quad (5.6)$$

These can be scaled relative to bulk concentrations and substituted into the Nernst-Planck equation to give the *Poisson-Boltzman equation*.² The specifics of this step, however, vary with geometry and ionic makeup and so are not presented in a generalized form. The steady-state case of a planar electrode and binary, symmetric electrolyte can be solved analytically to give the *Guoy-Chapman solution* [91]:

$$\frac{e\phi_{(z)-bulk}}{k_B T} = \ln \left[\frac{1 + \gamma e^{-z/\lambda_D}}{1 - \gamma e^{-z/\lambda_D}} \right] \quad (5.7)$$

where

$$\gamma \equiv \tanh \left(\frac{e\phi_{electrode-bulk}}{4k_B T} \right) \quad (5.8)$$

and the Debye length, λ_D , is defined as

$$\lambda_D \equiv \sqrt{\frac{\epsilon_0 \epsilon_r k_B T}{2e^2 \rho_0}} \quad (5.9)$$

where ρ_0 is the average density of ions in the bulk.³ Analytic solutions have also been found for the time-evolution of these equations in some relatively simple cases [55].

Numerically, this evolution can be simulated by discretizing the space, linearizing the equations, and solving the equations simultaneously at each incremental time (see Chapter 8). Though not working on CDI, Morrow performed such a simulation for a similar physical geometry [91]. These simulation results were then experimentally verified by Han [61]. Both the analytical solution and the Morrow and Han work, however, assumes a planar electrode rather than a porous one. This is a substantive difference and, thus, requires amendment before being applied to CDI.

² Depending on academic field, this may be alternatively referred to as part of Debye-Hückel Theory or Guoy-Chapman Theory.

³ If the model includes a Stern voltage, a voltage drop between the electrode and the nearest layer in the electrolyte, then $\phi_{electrode-bulk}$ must be reduced by that amount. At this point, it becomes a Guoy-Chapman-Stern model.

Integrating CDI's Porous-Electrodes

A CDI-specific model was proposed by Yang, et al., in 2001 [146], [147], which used the GCS theory (see Chapter 4), accounted for the high specific surface area of porous electrodes, but discounted the capacity of small pores where double-layers would overlap. Since then, however, Bazant, Thornton and Ajdari's weighty article presented a unified, mean-field theory for dynamics of locally flat electrode charging [12]. The theory describes two regimes: a "weakly nonlinear" regime for small voltages *or* early times, and "strongly nonlinear" regime for large applied voltages *and* long times.

Working with Bazant, Biesheuvel applied this to create a similar unified mean-field theory addressing dynamics in porous electrodes like those used in CDI [16]. This theory was then applied with Guoy-Chapman-Stern double layers to dilute, symmetric binary electrolyte case and solved numerically. From this full solution, they identify two analogous dynamic regimes, one in which the electrode behaves like a charging superconductor, and another dominated by slower, diffusion-limited desalination.

Though quite complete, this model cannot be used to determine the timescales of interest for PCCDI for several reasons. First, The model considers the electrode only and must be coupled to the bulk electrolyte to describe a full cell. Where it has been so coupled, they treat only the constant-voltage case. Further, they assume a stagnant diffusion layer at the electrode-bulk interface that may be a poor assumption with rapidly-varying potentials of pulsing. Finally, the model does not consider surface conductance.^{4,5} Among published work, however, this model is the most applicable to our goals. We therefore take this and the somewhat simpler work of Morrow and Han as "jumping-off" points for both theory (below) and simulations (Chapter 8).

Donnan Equilibrium Models for CDI

Despite limited application to our work, the *Donnan models* merit inclusion here. The first Donnan model divides pore volume into macropore and micropore volumes. Micropores are those for which we expect GCS solutions to overlap in the pore, with diameters $\sim 2\text{-}3$ nm ($\sim \lambda_D$ for a planar electrode case with the same electrolyte). Macropores are larger pores, typically up to $\sim 10\mu\text{m}$. Macropores are considered ion transport pathways, but not sorption sites; all surface area is considered to border micropores. The core of the micropore and all present ions are considered at the electrode potential. This micropore potential

⁴ This is equivalent to limiting application to small Dukhin numbers. I believe this will be a reasonable assumption in parts of the electrode near the ionic front, but not in regions with well-developed internal double layers such as late-timescale pores near the interface.

⁵ The model is also described as valid in the limit of thin-double layers. If this is a reference to the earlier assumption in Bazant, then it applies relative to gap distances. However, they state that this is relative to pore diameter, which seems less defensible in treating material with nm-scale pores.

differs from the macropore potential by a *Donnan voltage*, $\Delta\phi_D$. Ion i concentrations in macropores, $c_{i,ma}$, vs. micropores, $c_{i,mi}$ vary by a factor

$$c_{mi} = c_{ma} e^{-\left(\frac{z_i e \Delta\phi_D}{k_B T}\right)} \quad (5.10)$$

consistent with Boltzmann statistics, where $z_i e$ is ion charge, k_B is the Boltzmann constant, and T is temperature.

The *modified Donnan model* adds a Stern layer on the micropore surfaces, with the electrode and micropore potentials differing by a *Stern voltage*, $\Delta\phi_{St}$ [19]. Most recently, Biesheuvel has presented the *ampheric Donnan model*, which allows for charged functional groups on electrode surfaces [17], and is thus applicable to enhanced CDI, inverted CDI, and similar variations.

In contrast to Morrow’s simulation, the Donnan model includes porous media, but considers only electrophoretic ion motion into the bulk, treating diffusion as negligible. The model also simplifies pore-size distributions to the two classes of pores. This may be a particularly good approximation when dealing with hierarchical carbon aerogel monoliths, designed with pores of two distinct scales. Most relevant for our considerations, however, the model addresses only the equilibrium state. Though these models appear to provide good predictions for CDI performance, they cannot illuminate the timescales of interest for PC-CDI.

5.4 Our Charge-Current Model

Our theoretical work aimed to develop a model with both can be experimentally verified and can provide for interpretation of experimental results. More specifically, we sought equations to relate largely inaccessible variables of interest (time evolutions of ionic current and ionic charge in the double-layers) in terms of variables that can be measured experimentally (external current and applied potential).

In this section, we present the derivation of our *charge-current model*. Specifically, we describe the approach, define terms and variables, state assumptions, and present the mathematical derivation of the equation that may illuminate these otherwise inaccessible variables. We then derive a second equation that can be used to validate the theory, and discuss the application and limitations of this result.

Approach

Our approach has two steps. The first uses a physics-based paradigm, beginning with Poisson’s equation, to generate voltage equations and equivalent circuits for the CDI cell.

The second step involves confirming those results with circuit analysis based on *Kirchoff's Voltage law*. Neither the physics-based approach (see Morrow [91]) nor the use of equivalent circuits (see, for example, Ch 10 Bard and Faulkner [8]) is new. To my knowledge, however, this combination, this representation of a CDI cell, and this specific equivalent circuit are novel.⁶

Definitions

Our physical model consists of parallel, planar electrodes a finite distance apart, with substantially larger orthogonal dimensions (see Figure 5.1). For the physical model, we define the following terms:

- The gap distance, d_z , is the distance between the inner surfaces of the electrode.
- Orthogonal lengths, d_x and d_y , are the dimensions of the electrodes contributing to cross-sectional area.
- The Debye length, λ_D , is taken in z , and represents the scale of the electric double-layer.

In our prototype (see Chapter 6), $d_z \sim 1\text{mm}$, and d_x and d_y are 75 and 225 mm.

For the circuit analysis, we needed a circuit element where electronic charge in an electrode could be matched by electronic charge in the opposite plate, by ionic charge in the adjacent double layer, or by some combination of the two. We first envisioned a modified capacitor, where the external plates represent the electrodes, but where *part of* this capacitor has internal plates representing the double layers, separated by a resistance representing the bulk fluid (see Figure 5.2.a.). Unfortunately, traditional circuit laws cannot be easily applied to such a hybrid circuit element. This was resolved by replacing the hybrid element with an equivalent, albeit less intuitive, circuit composed of traditional elements (see Figure 5.2.b.).

For the circuit analysis, we define the following:

- q_e is the electronic charge on the positive electrode, the exterior side of capacitor C_2 .
- $-q_i$ is the net ionic charge in the electrical double-layer of the positive electrode, represented by the interior side of capacitor C_2 . Note that these charges are not differentiated between those lying in the Stern layers and those in the diffuse layers.

⁶ I am aware that much of this method may appear counterintuitive. This approach resulted from numerous iterative versions, each of which represented a small, and likely more intuitive, improvement on the last. I have not discovered any way to present this intuitively without presenting each of the previous, and ultimately flawed, iterations.




Figure 5.1: Our theoretical model of the CDI cell. Two parallel, planar electrodes (tan) a finite distance d_z apart, with substantially larger orthogonal dimensions. When the cell is externally raised to a constant voltage, electrons flow to and from the electrodes. In our model, this electronic charge lies in planes at the interior edges of the electrodes. Ions move toward the electrodes, attracted by this charge. In our model, (net) ionic charge is confined to two planes, each one Debye length away from each interface. Relative electric potentials are then defined by five linear regions. To make the simplifications explicit, the potential curve can be compared with those in Figure 4.9.

- Rather than a physical charge, $-q_c$ is the sum of q_e and $-q_i$, the electrical charge in the electrode matrix unmatched by ionic charge in the double layer, defined by

$$q_c \equiv q_e - q_i \quad (5.11)$$

In the circuit diagram, q_c is the magnitude of equal and opposite charges at the same potential, one that equalizes C_2 and the other that constitutes the charge on C_1 .

- i_e , i_c , and i_i are defined as the net currents flowing into C_2 , into C_1 , and through R_2 , respectively.
- The capacitance of C_1 is defined as

$$C_1 \equiv \frac{q_c}{\phi_{c-d}} \quad (5.12)$$

for reasons made explicit below and is here named *parallel-plate capacitance*.

- The capacitance of C_2 is defined as

$$C_2 \equiv \frac{q_e}{\phi_{b-c}} \quad (5.13)$$




Figure 5.2: My circuit models of the CDI cell. The first model (image a.) represents the CDI cell as a modified capacitor. The outermost vertical plates represent the CDI electrodes. Charge on the top half of those plates represent electronic charge in the electrode balanced by electronic charge in the opposite electrode, a parallel-plate capacitor with capacitance C'_1 . Charge on the bottom half of the electrodes represents electronic charge in the electrodes balanced by ionic charge in the adjacent double-layer, supercapacitors with capacitances C'_2 and C'_3 . As the ions can move through the bulk, they are connected via resistor with resistance R'_2 . R'_1 represents lead resistances. In the second, equivalent model (image b.), the CDI cell is represented by two capacitors in series with a resistor and capacitor in parallel. The electronic charge in the electrodes is q_e and the ionic charge in the double layer is $-q_i$. $-q_c$ is the shortfall between the two, the electronic charge un-matched by ionic charge. As the inner plates of the outer capacitors and the plates of the inner capacitors exist in the same location, $-q_c$ is a non-physical charge, always canceled by the q_c on the inner capacitor.

for reasons also made explicit below and is here named the *double layer supercapacitance*.

- R_1 is the sum of all resistance(s) in the leads.
- R_2 is the resistance of the electrolyte to ionic current.

Assumptions

The derivation employs several assumptions:

- The Debye length is much smaller than the electrode gap. This is often called the *thin double-layer approximation* [12]:

$$\lambda_D \ll d_z \quad (5.14)$$

- The electrode gap is much smaller than each of the orthogonal dimensions:

$$\begin{aligned} d_z &\ll d_x \\ d_z &\ll d_y \end{aligned} \quad (5.15)$$

- Any external electric fields are negligible, or equivalently, electric potential is constant outside of the cell:

$$\text{for all values of } z \text{ outside the cell, } \frac{\partial \phi}{\partial z} = 0 \quad (5.16)$$

This also implies that $\vec{E} = 0$. It is worth noting that assuming this for both sides of the cell also implies zero net charge within the cell, which will be consistent with the treatment.

- The double-layer supercapacitance is much greater than the parallel-plate capacitance:

$$C_1 \ll C_2 \quad (5.17)$$

- We are also treating the electrodes as simple planes, when in reality, they are porous surfaces with substantively more complex geometry. This will later prove to be a limiting factor for this theory.
- We assume negligible resistances across each electrode.
- We assume no Faradaic currents. This is embedded in the representation of the electrical double layers as capacitors C_2 and C_3 .
- We assume electro-neutrality in the bulk, outside of electrical double layer.
- We assume symmetric electrodes. Although not true for many CDI arrangements, this is consistent with our prototype.
- We assume a symmetric binary electrolyte, meaning equal and opposite charges and equal diffusivities. The experimental work was done with KCl, which comes close to this.

We note that the above two symmetry assumptions above imply equivalent response to electric fields and symmetric double layer capacitances, or

$$C_2 = C_3 \quad (5.18)$$

- Finally, we assume resistance and capacitances are, for each run, constant:

$$\frac{dR_1}{dt} = \frac{dR_2}{dt} = \frac{dC_1}{dt} = \frac{dC_2}{dt} = 0 \quad (5.19)$$

This and the assumption of planar electrodes are perhaps the most substantive assumptions and have a number of consequences discussed following the derivation.

Derivation of Charge-Current Equation

We begin with a simplified application of Poisson-Nernst-Planck approach. With 5.15, we can treat the space as uniform in x and y , but varying in z , effectively treating edge effects as negligible. With this simplification, Poisson's equation, 5.4, becomes:

$$-\frac{\partial^2 \phi}{\partial z^2} = \frac{\rho_q}{\epsilon_0 \epsilon_r} \quad (5.20)$$

In this model, charge is treated as planes of charge density $\rho(z)$, uniform in x and y .

With no external electric fields (5.16), the charge density, $\rho_q(z)$, dictates the shape (though not the absolute value) of the electric potential, $\phi(z)$, (and completely dictates the electric field, $E_z(z)$). The electric field is dependent only on the total of those planes at smaller (or larger) z values. We can therefore make the simplification that nonzero charge exists in four planes, at the two electrode surfaces and at two ionic layers at λ_D distant without affecting values for the electric field in electroneutral regions.

The double-layer voltage is then, with, equation 5.13,

$$\phi_{b-c} = \lambda_D \frac{q_e}{d_y d_z \epsilon_0 \epsilon_r} = \frac{q_e}{C_2} \quad (5.21)$$

By a similar method, voltage across the bulk (between the double layers) is then, with 5.11 and 5.12

$$\phi_{c-d} = (d_z - 2\lambda_D) \frac{(q_e - q_i)}{d_y d_z \epsilon_0 \epsilon_r} = \frac{q_c}{C_1} \quad (5.22)$$

The voltage drop across the bulk also induces the ionic current in the electrolyte:

$$\phi_{c-d} = i_i R_2 = i_i \frac{(d_z - 2\lambda_D)}{\kappa_b d_y d_x} \quad (5.23)$$

where κ_b is the ionic conductivity of the bulk. Combining 5.22 and 5.23, gives

$$q_c = i_i \tau \quad (5.24)$$

where τ is a timescale defined by

$$\tau \equiv \frac{\epsilon_0 \epsilon_r}{\kappa_b} \quad (5.25)$$

This treatment is consistent with traditional formulation of resistance as geometry-scaled conductivity and capacitance as geometry-scaled dielectric constants. The time constant, τ is then equivalent to an *RC time constant*, simply $R_2 C_1$ in more fundamental terms. With *Kirchoff's current law*, we can then build a differential equation that puts the observable i_e in terms of the unmatched charge, q_c ,⁷ via

$$i_e = i_c + i_i = \frac{dq_c}{dt} + \frac{q_c}{\tau} \quad (5.26)$$

To solve this differential equation, we multiply all terms by a function $\mu(t)$, with the crucial property that

$$\frac{1}{\tau} \mu(t) = \frac{d\mu(t)}{dt} \quad (5.27)$$

which can be found by rearranging,

$$\frac{1}{\tau} = \frac{\frac{d\mu(t)}{dt}}{\mu(t)} = \frac{d}{dt} \ln(\mu(t)) \quad (5.28)$$

integrating from a generic t_1 to a generic t_2 ,

$$\int_{t_1}^{t_2} \frac{1}{\tau} dt = \int_{t_1}^{t_2} \frac{d \ln(\mu(t))}{dt} dt \quad (5.29)$$

evaluating,

$$\frac{t_2 - t_1}{\tau} = \ln(\mu(t_2)) - \ln(\mu(t_1)) \quad (5.30)$$

and rearranging,

$$\mu(t_2) = \mu(t_1) e^{\left(\frac{t_2 - t_1}{\tau}\right)} \quad (5.31)$$

Returning to our differential equation, 5.26, we multiply by $\mu(t)$ and rearrange, now making the time dependencies explicit

$$\mu(t) \frac{dq_c(t)}{dt} + \mu(t) \frac{q_c(t)}{\tau} = \mu(t) i_e(t) \quad (5.32)$$

By our crucial property of $\mu(t)$, described in 5.27,

$$\mu(t) \frac{dq_c(t)}{dt} + q_c(t) \frac{d\mu(t)}{dt} = \mu(t) i_e(t) \quad (5.33)$$

⁷ Alternatively, the same two equations could be used to put i_e in terms of the ionic current, i_i , but the resulting differential equation is more complex.

And then, by the *product rule*

$$\frac{d}{dt}[\mu(t)q_c(t)] = \mu(t)i_e(t) \quad (5.34)$$

This can then be integrated (again, from generic t_1 to generic t_2)

$$\int_{t_1}^{t_2} \frac{d}{dt}[\mu(t)q_c(t)]dt = \int_{t_1}^{t_2} \mu(t)i_e(t)dt \quad (5.35)$$

Evaluating the left-hand-side and solving for $q_c(t_2)$,

$$q_c(t_2) = \frac{\mu(t_1)q_c(t_1) + \int_{t_1}^{t_2} \mu(t)i_e(t)dt}{\mu(t_2)} \quad (5.36)$$

And expanding $\mu(t)$ and $\mu(t_2)$ using 5.31,

$$q_c(t_2) = \frac{\mu(t_1)q_c(t_1) + \int_{t_1}^{t_2} \mu(t_1)e^{\left(\frac{t-t_1}{\tau}\right)}i_e(t)dt}{\mu(t_1)e^{\left(\frac{t_2-t_1}{\tau}\right)}} \quad (5.37)$$

and simplifying, becomes what we will call the *charge-current equation*:

$$\mathbf{q_c(t_2) = q_c(t_1)e^{\left(\frac{t_1-t_2}{\tau}\right)} + \int_{t_1}^{t_2} i_e(t)e^{\left(\frac{t-t_2}{\tau}\right)} dt} \quad (5.38)$$

$$\tau \equiv \frac{\epsilon_0 \epsilon_r}{\kappa_b}$$

This equation puts the unanswered charge, q_c at some time t_2 in terms of a known q_c at a previous time t_1 and the external current function, $i_e(t)$, from t_1 to t_2 . In practice, $q_c(t_1)$ can be reduced to zero, eliminating the first term. It is worth stating that q_c is proportional to electric field across the bulk.

With $q_c(t)$ known, $q_i(t)$ can be determined by several methods, but most easily by integrating 5.24, giving

$$q_i(t_2) = q_i(t_1) + \frac{1}{\tau} \int_{t_1}^{t_2} q_c(t)dt \quad (5.39)$$

This then describes the growth of the electric double layer.

Conceptually, the mathematical result can be understood as describing each incremental bit of charge, $i_e dt$, entering the cell as entirely unanswered charge, q_c . The contribution of that bit to the total unanswered charge then decays exponentially with time constant τ from the time it entered the cell (see Figure 5.3).




Figure 5.3: Interpretation of the charge-current result. At left, the external current, i_e can be seen as charging both the external and internal capacitors, increasing both q_e and q_c . Superimposed on this current, the ionic current, i_i discharges C_1 through resistor R_2 . At right, an arbitrary charging curve. Each incremental bit of the externally provided charge, $dq_e = i_e dt$, initially contributes to q_c (in blue). Each contribution, however, decays with a time constant $\tau = R_2 C_1$ from the moment of arrival. The red represents the charge that has been transferred via ions, q_i .

Confirmation Method

This theory also predicts the voltage drops across each element, based on the charges determined by 5.38. As the total cell voltage, ϕ_{b-e} , can be measured in lab, this gives a method to experimentally confirm the charges q_c and q_i . Specifically, via *Kirchhoff's voltage law*,

$$\phi_{b-e} = \phi_{b-c} + \phi_{c-d} + \phi_{d-e} \tag{5.40}$$

With 5.12, 5.13, and 5.18, this becomes

$$\phi_{b-e}(t) = 2 \frac{q_e(t)}{C_2} + \frac{q_c(t)}{C_1} \tag{5.41}$$

To be valid, the externally measured cell voltage over time, $\phi_{b-e}(t)$, must be a consistently weighted combination of the temporal functions $q_e(t)$ and $q_c(t)$.

Limitations-Effects

In many ways, our assumptions reduce the electrical double-layers to the Helmholtz model. Non-net-neutral ionic charge is assumed to exist in a single infinitesimally small Helmholtz

plane of infinite concentration. At short timescales, this proves reasonable, but diffusion from this plane is not fully considered. Neglecting diffusion ignores effects that stem from both diffusion back into the bulk and diffusion into the electrode itself.

The first diffusion effect moves the location of counter-ions in the double layer from the Helmholtz planes to \sim the Debye length. This means the strong electric field—and corresponding steep slope of the electric potential—persist over a larger distance, thus increasing the total voltage drop over the double layer. Though we did not include diffusion, this effect is tacitly included in the charge-current model, to first-order, by placing the counter-ion mono-layer at the Debye length, λ_D , rather than at the plane of closest approach for either saturated or unsaturated ions (as, say, in the Helmholtz model).

Ion diffusion into the electrode bulk also affects ionic conductivity. In our model, ionic conductivity, κ_b , is treated as constant (via 5.19) and uniform (a result of several assumptions). While a constant κ_b is a reasonable assumption in many parts of the cell, diffusion from the ion planes to the bulk will resist continued ionic current in the adjacent areas, lowering effective conductivity.

Uniform ionic conductivity also breaks down in nonlinear regimes.⁸ Early on, or with low voltages, the attraction of counter-ions and repulsion of co-ions are nearly symmetric. However, counter-ions can be added nearly indefinitely, but there are a finite number of co-ions to expel. As the response continues, therefore, co-ions will be expelled from increasingly distant planes, beyond the Debye length. As this region cannot support large discrepancies in co- and counter-ion concentrations, the counter-ion concentration is also suppressed here. This creates a region of lower total ionic conductivity (see Figure 5.4). This effect is theoretically treated in Bazant [12] and Biesheuvel [16], shown experimentally by Suss [131], and demonstrated through simulation in Chapter 8 of this work.

Finally, and likely most significant, the electrodes are assumed to be planar, when in practical CDI experiments, they are porous. The model, therefore, does not account for ion motion into the electrode pores. This motion, beginning as ions arrive, gradually changes the surface area of the double-layer from the *projected area*, $d_x d_y$, to the much larger *tortuous surface area*. This effect, therefore, means C_2 is not constant as asserted in 5.19.

All three drawbacks to the charge-current model depend on diffusion or diffusion-like effects. Diffusion, in turn, scales with concentration gradients. As a result, this theory proves internally consistent at early timescales—where diffusion into the electrode and bulk are negligible. However, it fails its confirmation step at later timescales—where diffusive

⁸ Response has been divided into a linear (low voltage and early timescales) and non-linear regime (larger voltage or late timescales, or, I believe, low ionic strength) [55]. Bazant [12] and Biesheuvel [16] further distinguish between the weakly nonlinear and strongly nonlinear regimes.




Figure 5.4: Spatially local decreases in overall ion concentration near the double-layer. Time, $t = 0.5\tau_D$ (solid), $1\tau_D$ (dot), $2\tau_D$ (dash), $4\tau_D$ (dot-dash), $20\tau_D$ (long dash), where τ_D is *Debye time*. Image is from Bazant [12].

fluxes are significant and double-layer areas have increased.⁹ I therefore limit application to experimental results to those early timescales for which it passes confirmation steps. For this application, see Section 7.3.

As an aside, the confirmation method offers an interesting possibility. If this theory could be epistemologically verified by another method, then the voltage equation could be solved for a non-constant C_2 . This would illuminate the evolving double-layer area as ions diffuse into the bulk. As it stands, however, there would be no way to independently verify that interpretation. This approach would also require that the (likely smaller) effects of diffusion into the bulk be either separately incorporated or neglected.

⁹ This theory is also expected to fail for exceptionally fast timescales following an instantaneous charge, as the water must first respond to the electric field. During this time, ϵ_r , and thus C_1 , are non-constant. That said, charging the cell fast enough to observe this would be technically challenging.

5.5 A Case for Numerical Simulation of Late-Timescale Effects

While the charge-current equation provides some insight, without incorporating bulk diffusion, local concentration suppressions, and movement into the electrode, it cannot generate realistic estimates for total capacitance, decay of bulk field, or other longer-timescale effects. In this section, we explore what would be required to capture bulk effects and ion infiltration of the electrode, finally arguing for numerical simulation rather than analytic treatment.

Requirements to Capture Late-timescale Phenomena

Capturing bulk effects requires including the diffusion term in the Nernst-Planck equation, 5.5, which in turn, requires spatially and temporally varying concentration functions for each species of ion. This challenge is equivalent to solving the Poisson-Nernst-Planck equations in the nonlinear regime(s).

Capturing transit into the electrode and resultant effects also adds complexity. The conductive electrode will tend to neutralize electric fields in its interior. However, such fields may still be a factor, near the interface, in larger pores, and after ions have entered the pore. Where diffusion dominates, it will be governed by a reduced effective diffusivity, D_e . This would vary with material properties like porosity and tortuosity, but also have contributions from both movement through the pores and movement along the surfaces, which, in turn, vary spatially and temporally.

The Case for Simulation

The complexity required to capture these phenomena suggests an analytical approach would be perilous. Even with planar electrodes, capturing diffusion effects within the bulk requires solving the Poisson-Nernst-Planck equations in the nonlinear regime(s). While it may be possible to solve these equations analytically, as has been done for the linear regime [55], this would still not include effects due to the porous electrode. Incorporating the porous electrode means that in addition to spatially and temporally varying potential and concentration functions, we must link similarly-varying scaling functions for conductivity and electrophoresis. Consequentially, we have not pursued an analytic approach for capturing late-timescale behavior.

Instead, we argue that the complexity of capturing these effects recommends a numerical approach. Chapter 8 shows our work using numerical simulation to generate more complete models. This can be seen as either an alteration of Morrow [91] and Han [61] that includes porous electrodes, or a variation on Biesheuvel and Bazant's work [16] that treats bulk and

electrodes simultaneously.¹⁰ Such simulations can confirm near-term timescales given by the charge-current model—such as voltage collapse—and generate experimentally verifiable predictions longer-term timescales such as complete screening of bulk field.

¹⁰ In addition to the combined treatment of bulk and electrolyte, there are other substantive, but nuanced differences in the assumptions between this effort and that of Biesheuvel.

Chapter 6

Pulse-Charged CDI Prototype Design

6.1 Introduction

In Chapter 4, we described the pulse-charged CDI arrangements for increasing effective ion mobility, and in Chapter 5, we presented theory to be used in experimental work. In this chapter, we will describe the iterative design, fabrication and testing for the prototype cell that will become the testbed for the experimental work covered in Chapter 7.

We developed our CDI prototype to achieve two goals. First, we sought to demonstrate traditional CDI operation, replicating published performance, both to confirm our understanding of the technology and create a reproducible baseline. Second, we wanted a device capable of operating in pulsed-CDI modes,¹ such that we could perform side-by-side comparisons.

In this chapter, we describe the design requirements for such a PC-CDI prototype and present the initial design, based on work by the Farmer group. We then chronicle many of the prototype failures, diagnoses, and design iterations that resolved them. Finally, we present the functional design used for much of the experimental work along with the first multi-cycle desalination results.

6.2 Design Requirements and Initial Design

In this section, we present design requirements for our CDI prototype, divided into requirements for the CDI cell and requirements for the electrode material. We then describe the initial design concept, based on work by the Farmer group, which was used as a starting point.

¹ The prototype work was begun when voltage-pulsed CDI and the Wien effect were still under consideration. As the kV-range concepts were dropped, the goal focused on pulse-charged CDI, and requirements relaxed substantively.

CDI Cell Requirements

There were a number of considerations for the design of the CDI cell. First, as we aim to affect and observe increases in effective ion drift rates, we want to maintain steady, uniform, and reproducible flow patterns within the cell. For this reason, We chose not to include a non-conductive, porous spacer material, which is sometimes used to keep electrodes separate, but which likely creates more nuanced flow patterns. We can also expect increases in ion drift rates to be more reliably observable if water flow remains in the laminar regime. For a traditional CDI cell, the Reynolds number can be taken as

$$R_e \approx \frac{ud_z}{\nu} \quad (6.1)$$

where u is the magnitude of fluid velocity (here taken as volumetric flow per cross-sectional area), d_z is distance between the electrodes, and ν is the kinematic viscosity of water. For flow rates of ~ 20 mm/s and a gap ~ 1 mm, $R_e \sim 20$, well within the laminar flow regime.

As ion drift rate is a strong factor in minimizing hydraulic residence time, any attempt to increase flow rate per electrode material must also consider the ion travel distance. In a traditional CDI cell, this distance is the gap distance between electrodes, perpendicular to flow rate. Broadly speaking, this should be minimized, but without insulating spacer material, risk of electronic shorts must also be considered.

Typically, traditional CDI cells have relatively low pressure requirements. As a result, pumping does not contribute substantively to per-volume energy requirements. Indeed, this is the principle advantage of CDI over reverse osmosis when treating brackish water. This constitutes another competing consideration at very small gap distances (see, for example, Suss [131]).

We also wanted a large projected surface area for the electrodes. As a practical matter, this increases both total surface area and cell capacity, thus increasing the time between recharge phases. It also allows for larger—and hence more easily detectable—removal rates. Further, much of my theory rests on the assumption that cross-sectional electrode dimensions are much larger than the electrode gap. Therefore, we used a somewhat larger ($\sim 3x$) unit electrode than designs upon which this was based.

In a related consideration, we wanted the cell to be scalable. If we observed a potential phenomena, but did not have sufficient removal to resolve the effect, a scalable device would allow us to increase the effective surface area of the cell in a subsequent run. Alternatively, if we observed advantageous effects, it would be worthwhile to ensure that they persist at more practical scales.

CDI Electrode Requirements

There are also a number of materials requirements for an effective CDI electrode [105]. Specific surface area should be high, to allow larger capacity and longer removal phases. High electrical conductivity is also desirable. The material should be mechanically stable in the presence of fluid flows. Finally, the materials should be electrochemically stable over a wide variety of possible conditions.

Our specific goals provide additional electrode requirements. As we seek to develop a baseline comparable to published results, the electrode material should be selected from those widely-used and well-characterized. As we aim to increase effluent flow per capital cost, the electrode should be among the most affordable that meet the above requirement. These two requirements effectively precludes many of the more recent and exotic materials being explored.

Initial Design

Considering these requirements, we based the initial prototype concept on work done by the Farmer group (see [43], [44], [47], and especially [41], [38]). In this design, water flows upward through a serpentine path (see Figure 6.1). On each lateral pass, it flows between oppositely charged carbon electrodes attached to metallic electrode plates which act as structural support and current collectors. The plates are separated by electrically insulative spacers and sealed with compressible gaskets. At the top and bottom are mechanically robust header plates. The spaces between the outer electrode plates and header plates create initial and final passes that help keep fluid flow uniform. The stack is aligned and compressed with threaded rods, and can be scaled by adding additional electrode layers.

The header plates, electrode plates, and initial gaskets were cut from stock material using a water jet. The original header plates were made of 1/4 inch 316 stainless steel plate;² the original electrode plates were from 1/8 inch 6061 T6 aluminum sheet.³ The gaskets were cut from 1/16 inch black medium-strength (60A Durometer) EPDM rubber.⁴ White delrin spacers, nylon #10-24 threaded rods, #10 washers, #10 lock-washers, and #10-24 nuts were all commercially available parts. Electrodes were attached to the electrode plates with silver-impregnated conductive epoxy (Tiga Silver 920H).

The carbon nanofoam electrode sheets were cut to just under 75 x 225 mm (3 x 9 inches). The 1.52 mm spacers between electrode plates, with the electrode paper thickness of ~ 0.25

² Type 316 stainless steel is frequently used in seawater environments, and due to high molybdenum content, has improved resistance to oxidizaion in such environments.

³ McMaster-Carr #89015K28, here included as we later present phenomena regarding this material.

⁴ McMaster-Carr #8610K72, also included as we later present phenomena regarding this material.




Figure 6.1: At left, a conceptual diagram of our initial prototype design, based on work by the Farmer group. Water follows a serpentine path upward through the stack. After each pass, it passes through a gap in the electrode plate to the next level. The stack can be scaled to include any number of electrode plate pairs. At upper right, an isometric projection of the actual electrode plate design. The protruding tab is for electrical connection. At lower right, a photograph of the cell being assembled. The metallic material is the electrode plate, and the black material is the carbon aerogel electrode, which has been affixed to the plate with conductive epoxy. Also visible are 20 nylon threaded rods used to compress the stack.

mm, gives an estimated electrode-electrode gap of 1.0 mm, although epoxy thickness may have had an effect in early versions.

6.3 Failures, Diagnoses, and Iterative Design

In this section, we discuss three major failures of the initial design, their respective diagnoses, and the design changes that resulted. We then briefly describe many of the smaller challenges and improvements. As these issues occurred with overlapping time-frames, they are not presented chronologically. We have instead endeavored to disentangle the issues and present each one independently.

Chasing Leaks

Assembly of the initial prototype appeared to go well. When we pumped water into it, however, nearly every gasket connection appeared to be leaking.⁵ We tried to tighten the nuts on the threaded rods, compressing the stack and the gaskets, but soon reached the limit of the nylon as the threads began to deform. We then attempted to resolve the issue by increasing compression with several large c-clamps. This allowed us to further tighten the stack, but the problem persisted.

In the end, there were three factors, each of which were likely sufficient to prevent a good seal. First, the weak nylon rods, chosen because they were electrically insulative, were not capable of providing sufficient compression force even when the force was distributed over 20 rods. Second, the initial gaskets (see Figure 6.2), especially those adjacent to the header plates, had large contiguous areas of material. This severely limited their compressibility, and left the areas where the seal was crucial effectively uncompressed. Third, the delrin spacers, while specified to be 0.06 inches (1.52 mm) thick, were actually 1/16 inches (0.0625 inches or 1.59 mm) thick, and they did not allow for sufficient compression.

The three-part solution involved replacing the nylon rods with stainless steel ones, trimming the gaskets to a thin width around the water-containing area, and turning the 76 spacers to the smaller required thickness on a lathe. Since the new stainless steel rods were electrically conductive, however, they threatened to short the electrodes. Therefore, we also bored out the holes in the electrode plates, and replaced the spacers with shoulder-spacers, which were press-fit into the electrode plates, acting as an insulator between the rods and the electrode plates. The shoulder-spacers also had to be turned down to reach the 1.52 mm goal.

Chasing Shorts

Though there were many shorts over the course of design and testing, we here present the most counterintuitive—and therefore most interesting—case. When we began to apply a voltage to the electrodes, we observed a nontrivial background current.⁶ Even without a conductive electrolyte, there was some electrical pathway between the electrodes. Measuring the resistance between the electrodes of the drained cell with a multimeter gave a resistance $\sim 80 \Omega$ and gradually falling.

We tested the resistance of all materials that could be making contact: the delrin spacers, the black EPDM gaskets. They all tested as insulative. We monitored the resistance as we

⁵ In the labbook, see work from 13 May 2014, run 1 to 14 July 2014, run 3.

⁶ In the labbook, see 8 Dec 2014 for awareness of the problem, 10 Dec 2014 for diagnosis, and 12 Feb 2015, run 9 for implemented resolution.




Figure 6.2: At left, a photo showing the original black EPDM electrode-electrode gaskets with large contiguous areas enclosing the nylon threaded rods. At right, a photo showing the thinner electrode-electrode gaskets, here in white EPDM. Other differences presented here include the move from 20 black nylon rods to 12 stainless steel ones and the change from spacers between every level to spacers spanning header plate to header plate.

disassembled the cell, but it rose to an open circuit before we got it apart. Imagining the electrodes were touching, either because they were lifting off the electrode plates or because the epoxy layer was thicker than we expected, we assembled a cell using bare electrode plates with no electrode material attached. The short returned. The answer came after repeated assemblies and disassemblies while monitoring the resistance.

The gaskets were the culprit. Black EPDM rubber is black because of the presence of carbon. At atmospheric pressures, this carbon is sufficiently discontinuous that the bulk material is electrically insulative. Under compression, however, the carbon becomes more contiguous, sufficient to grant the bulk material non-trivial conductivity. A cursory search suggests that this pressure-dependent resistivity in carbon (and carbon-bearing) polymers has been known since at least 1980 [45]. This effect was used (in reverse) by Bell Communications Research to design a fuel tank leak detector in 1989 [13], which won an R&D100 Award in 1991 [123]. To address this, we replaced all gaskets with white EPDM rubber, which is embedded with (electrically insulative) titanium oxide instead of (electrically conductive) carbon.

Chasing Chemical Reactions

Another interesting issue arose when we first began flow-through trials.⁷ Over a series of five runs, we intermittently observed bubbles in the effluent and visibly turbid effluent

⁷ In the labbook, see work from 13 June 2014, run 2 through 26 Aug 2014, run 4.

(see Figure 6.3). Neither was true of the influent water. At the end of one run, a funnel in the effluent line was coated in a fine whitish grit. In addition, effluent concentration rates had become increasingly unaffected by the applied voltages in these early runs. When we eventually opened the prototype cell, I observed:

whitish flakes on the interior edges of the electrode plate-electrode plate seal. I noticed the electrode coming off the electrode plate before I found a place where the epoxy was still attached firmly enough to cause the electrode to break when I tried to lift it. Underneath the electrode was... a liquid with white particulate visible within. I note that no such materials can be observed ‘upstream’ of the electrode pair.




Figure 6.3: Photographs from 21-26 Aug 2014. The upper left image shows turbid effluent observed during a run. The right image shows a large fraction of the black electrode that had become detached, as well as the white material visible underneath it. The silver epoxy lines are visible on the detached piece above. The lower left image shows the pitting visible on the aluminum electrode plate after sanding off the remaining epoxy and electrode material. The four horizontal lines where epoxy had been were visible in the pitting pattern.

Clearly, some undesirable chemical reaction was occurring, which was also degrading the attachment of the electrodes, and thereby the performance of the cell.

We considered two initial hypotheses. First, changing conditions had led to the precipitation of a compound involving K from the KCl in the influent: KCl , K_2CO_3 , or $KHCO_3$. Unfortunately, those are all relatively soluble. Alternatively, if Ag was being electrochemically freed from the epoxy in an effective anode, it could form the relatively insoluble $AgCl$. Indeed, basic CDI operation would concentrate Cl in that electrode anyway. The other half-reaction could then explain the gas bubbles. For example,



would have a minimum voltage requirement of 0.80 V at standard conditions. We had applied cell voltages as high as 1.4 V, and we expected some acidification in the pores of the negatively charged electrode.

As a diagnostic step and possible remedy, we affixed new electrodes to new electrode plates, using a carbon-based conductive epoxy (Duralco 127). Over the subsequent runs, however, we again observed poor removal and occasional bubbles in the effluent.⁸ When we again opened the prototype cell, we saw localized areas where the electrode had risen from the electrode plate, and white material had deposited or formed around the electrode periphery (see Figure 6.4). It seemed quite possible that this was the same phenomena, but observed at an earlier stage.




Figure 6.4: Photographs showing localized rising in the black electrode (left) and buildup of an initially unknown white substance near the electrode edges (both images). Photos from 21 Oct 2014.

While we still considered salt deposition and potential issues with the new epoxy, we eventually focused on the aluminum electrode plates. Eliminating these plates and fabricating new ones, however, would be a major step, to be avoided if unnecessary. As a preliminary

⁸ In the labbook, see 18 Sept 2014, run 5 and 15 Oct 2014, run 6.

diagnostic, therefore, we looked to Scanning Electron Microscope Energy-Dispersive X-Ray Spectroscopy (SEM-EDS). This method bombards a sample with electrons, varying the per-electron energies over ~ 0.1 -10 keV. Each element preferentially absorbs and disperses energy of particles at specific energies, corresponding to excitation states within their atoms. By seeing which energies are absorbed in the material, one can therefore determine which elements are present. We used SEM-EDS to examine samples of the both the white material and the electrode (see Figures 6.5, 6.6).




Figure 6.5: SEM-EDS results for one sample of the white buildup found at the electrode edges. At left, the SEM image. At right, spectroscopy results showing adsorption-dispersion at $K\alpha$ absorption lines for both oxygen and aluminum. There is also a small peak indicating the presence of carbon.

The scans showed the white substance was principally aluminum and oxygen. It seems likely that aluminum was oxidized into solution via



and formed relatively insoluble oxides (e.g. Al_2O_3), oxyhydroxides (e.g. $AlO(OH)$), or hydroxides (e.g. $Al(OH)_3$). Unfortunately, it is difficult to identify the exact species with SEM-EDS, as the technique has limited resolution with respect to relative contributions of elements and does not reliably resolve hydrogen. It is worth noting that the potential requirements for aluminum oxidization are substantially lower than those for silver oxidization. As the relative kinetics are unknown, however, the presence of the inferred aluminum reaction does not preclude the silver one; it is still possible they both occurred in the runs with the silver epoxy.




Figure 6.6: SEM-EDS results for one sample containing a particle of the white material on the nanofoam electrode. At upper left, the SEM image showing the white particle against the black electrode. At upper right, the spectrum shows peaks for carbon, oxygen, and aluminum $K\alpha$ lines. The lower images show the spatial locations of counts for carbon (left, green), oxygen (center, blue), and aluminum (right, red). The carbon locations are consistent with the electrode material in the background, but the oxygen and aluminum locations cluster around the particle.

To avoid oxidizing the aluminum electrode plates, we designed electrode plates made from 316 stainless steel. To reduce weight and cost, and because steel is stronger, we decreased the thickness from 1/8 inch to 1/32 inch. The thinner material, however, was incompatible with the shoulder-spacers, so we moved to a single spacer between header plates, passing through the electrode plates. In this variant, spacing between any two adjacent plates is maintained by the compressed gasket (see Figure 6.7).

Other Minor Iterations

Throughout the design process, we identified several issues leading to minor design changes. We briefly present them here.

- Initially, the edges of the electrodes were even with the edges of the header plates, making accidental physical contact easy. This would result in both a physical shock and a change in the electrode charge. Subsequent designs left the electrode plates inset from the header plates.⁹
- Initially, electrodes and gaskets had a small overlap, to limit the amount of bare electrode plate exposed to the solution. Unfortunately, compressing the gaskets shattered the electrodes. The interface was redesigned with a small gap between electrode and gasket.¹⁰
- The first attempt at using Duralco 127 carbon-based epoxy and stainless steel electrode plates led to poor adherence and eventual separation. This was resolved by abrading the electrode plate in the area where the epoxy would be used, but not in the gasket-interface area.¹¹
- In early versions, application by hand created uneven distributions of epoxy and thus, non-planar electrodes. This created non-uniform gap distance and risked creating shorts. This was resolved by smoothing the epoxy with a *doctor blade* (typically used to smooth inks) calibrated to create a 100 μm layer.¹²
- Finally, we changed the stainless steel header plates to transparent polycarbonate, which allowed for a lighter cell and better troubleshooting when observing bubbles and precipitates. Additionally, it reduces a parallel-plate capacitance that existed between the steel header plates, electrically connected to each other, and adjacent electrode plates.

6.4 Current Prototype Design and Proof-of-Concept

After many iterations, we achieved a stable design that could both perform as a traditional CDI cell and serve as the testbed for experimental investigation of the timescales necessary for pulse-charged operation. In this section, we present the elements of this design, formal drawings of the most nuanced custom parts, and our first, proof-of-concept results demonstrating traditional CDI operation.

⁹In the labbook, see 13 May 2014, run 1.

¹⁰In the labbook, see 20 May 2014, run 1.

¹¹In the labbook, see 23 Jan 2015.

¹²In the labbook, see 2, 23 Feb 2015




Figure 6.7: A conceptual cross-section of the current, stable design. Changes from the initial concept can be seen by comparison with Figure 6.1.

Current Design

The current design retains the serpentine fluid path of the original concept, but almost every other element has seen some alteration (see Figure 6.7). Structurally, transparent polycarbonate header plates (see Figure 6.11) are compressed with twelve stainless steel threaded rods. Delrin spacers define the total distance between header plates, but distribution of this space is defined by the compression of thin white EPDM gaskets. The electrodes are supported by 316 stainless steel electrode plates (see Figure 6.10), which also serve as current collectors. These plates are abraded before the electrodes are adhered with Duralco 127 carbon-based conductive epoxy. Delrin header blocks interface the slotted gap in the header plate to tubing used for influent and effluent. The threaded rods, washers, lockwashers, and nuts are all off-the-shelf components.

The electrodes are a Grade II carbon nanofoam paper produced by Ocellus, Inc. of Livermore, CA and procured through Marketech International, Inc. of Port Townsend, WA. The nanofoam paper has the following manufacturer-specified physical properties: density of 0.4-0.5 g/cm², specific surface area of ~600 m²/g, and specific capacitance of 28-30 F/g. The paper thickness is 0.010 inches (250 μm).

Design Drawings

At the end of this chapter, we include formal design drawings for the more nuanced parts. Specifically included are the electrode plate (Figure 6.10) and header plate (Figure 6.11).

Proof-of-Concept

The stable design was operated in traditional CDI arrangement, providing our first data on removal and regeneration over multiple cycles (see Figure 6.8).¹³ This was performed in a semi-batch mode, where effluent returns to the electrolyte reservoir. As a result, influent concentration is not constant and we report the difference between effluent and influent concentrations. We considered this a proof-of-concept, demonstrating that the prototype could perform as a capacitive deionization cell. These results can be compared with, for example, those of Gabelich [47], reproduced here in Figure 6.9.

While there have still been intermittent challenges and minor design adjustments, the basic design has remained fixed since this version. We have, therefore, taken this point as the transition between prototype development, which focused on creating a viable CDI testbed, and experimental work, which focused on understanding the timescales required for operation of pulse-charged CDI. This experimental work continues in Chapter 7.

¹³ In the labbook, see 18 Feb 2015, run 9.




Figure 6.8: Data from the first successful multi-cycle run with my CDI prototype. The run was conducted in a semi-batch mode, where effluent returned to the source reservoir. The data represents four cycles, each with a removal phase and a regeneration phase. The top image shows the difference between effluent and influent conductivities. Negative values represent removal; positive values represent regeneration. The middle image shows external current, and the lower image shows applied cell voltage. The abscissa for all charts is volume of water exiting the cell. There was a short break in the applied power at $\sim 700\text{mL}$ which seems to have created an uncharacteristic third cycle. The apparent overshoot in the regeneration phase is thought to be an artifact of the semi-batch method.




Figure 6.9: Image from Gabelich, et al., [47], showing typical data for traditional CDI cycles, included for comparison with Figure 6.8. It shows concentration and voltage for multiple complete cycles. The data was collected in a closed-loop configuration using 1 L, 0.004 M NaCl solutions at 1.0 V, 100 mL/min flow rate, and pH 7. When drawing comparisons, note our use of relative conductivity, $\kappa - \kappa_0$, and their use of fractional concentration c/c_o , as well as our use of volume and their use of time in the abscissa. We do not understand why c/c_o is reported consistently below 1, when we would expect it to rise above 1 during regeneration phases




Figure 6.10: Design drawing for CDI Electrode Plate, created in SolidWorks.




Figure 6.11: Design drawing for CDI Header Plate, created in SolidWorks.

Chapter 7

Pulse-Charged CDI Experimentation

7.1 Introduction

In Chapter 4, we presented the case for using capacitive deionization (CDI) for desalination of brackish drinking water in developing world communities. However, capital cost per throughput poses a potential barrier to application at scale in such settings. We saw that pulse-charged CDI (PC-CDI) has the potential to increase ion-drift, and therefore throughput, without inducing substantive redox reactions, and that even traditional CDI has not been optimized for maximum ion drift. Finally, we saw that the key to PC-CDI lies in understanding the relative timescales for pulse rise, shielding of the electric field in the bulk, redox reactions, voltage collapse (via internal equilibration), and decay from any intermediate energized states.

In Chapter 5, we presented theoretical framework that allows us to capture evolution of the electric field in the bulk, based on externally observable parameters, valid at early timescales. Specifically, we derived the charge-current equation (Equation 5.38). The framework includes a measure of internal consistency, so the validity can be confirmed at a given timescale. In Chapter 6, we detailed the development of our prototype CDI cell and experimental testbed.

Goals & Approach

This experimental work presented in this chapter¹ aims to identify the early timescales relevant for PC-CDI and their dependence on electrolyte concentration. In PC-CDI, cell charging takes place in two phases. In the *active-charging phase*, the cell is connected to a voltage source and electronic current flows to the electrodes. This is equivalent to the traditional CDI arrangement, but is maintained for a much shorter duration. In the

¹ The work reported here represents efforts from the somewhat misleadingly named *Impedance Waveforms* effort 14 Dec 2015 - 1 Apr 2016 (pages 93-119 in the labbook) and from the *Firm Pulse Waveforms* effort, 31 July 2016-15 Aug 2016 (pages 121-137 in the labbook).

relaxation phase that follows, the ions catch up. The cell is effectively isolated from charging circuitry, ions flow to the electrode double-layers, providing more complete shielding of the electric field, and the voltage drop across the bulk collapses.

We seek to understand the timescales within each phase. Specifically, we expect to resolve the timescale for initial ionic shielding of the electric field during the active-charging phase, as well as timescale(s) for the collapse of cell voltage in the relaxation phase. Characterizing timescales for the two phases was performed through two separate experimental series. Throughout the chapter, we will differentiate them by referring to the *active-charging-phase* (or *ACP*) experiments and the *relaxation-phase* (or *RP*) experiments.

In this chapter, we describe our experimental methods, including physical apparatus and laboratory procedures for each series. We next present our analysis, including corrections to current measurements, application of the charge-current equation, and verification steps. We then discuss results and for both ACP and RP data. Finally, we present our conclusions about the relevant timescales, and again make a case for simulation as a means to confirm and extend these results.

7.2 Experimental Methods

In this section, we describe the physical apparatus for our experiments, including ACP electronics, RP electronics, and electrolytes used. We then describe the procedures for ACP and RP experiments.

Experimental Apparatus

Data was collected under two experimental arrangements. The first active-charging-phase arrangement was used to evaluate timescales for the shielding of the electric field under a constant voltage bias. The relaxation-phase arrangement was used when measuring voltage collapse after the cell has been rapidly charged then isolated from the charging circuit. In both cases, the timescales were measured at multiple solution concentrations.

Active-Charging Phase Electronics In ACP measurements of initial ionic shielding timescales, voltage pulses were produced with a Stanford Research DG645 delay generator.² Current was captured with a Pearson 6027 current monitor. Both the pulse output and current monitor outputs were observed and recorded with a Tektronix MDO 3014 oscilloscope. This electronic arrangement is shown in Figure 7.1.³

² Preliminary data was taken with a Beckman Industrial FG2A function generator in place of the Stanford Research delay generator, but all data presented here used the DG645.

³ This electronic arrangement was recorded for various work 14 Dec 2015 - 13 Mar 2016 (pages 93, 95, 107, and 110 in the labbook).




Figure 7.1: Electronic arrangement for observing the initial ionic shielding of the electric field under constant voltage bias (ACP experiments). Green lines indicate the wires carrying the active signal. Orange lines represent (low current) wires used to observe voltages. Black lines are ground. The small box represents a BNC splitter.

Relaxation Phase Electronics In the RP effort, we sought to measure the timescale(s) for voltage collapse following a brief charging period, then isolation from the charging circuit. This required a circuit capable of both creating a rapid voltage rise and delivering large amounts of charge. To accomplish this, we adapted a circuit used in beam control of LBNL’s Advanced Light Source. A basic circuit diagram for this custom charging circuit is given in Figure 7.3. This circuit was powered by a TDK-Lambda Genesys 1U750W power supply. The pulses were triggered by a Stanford Research DG645 delay generator and supplied by a BK Precision 9150 power supply. Output was passed through a Fairchild 1N4447 diode. Current was captured with a Pearson 6027 current monitor. Trigger voltage, Pearson output, custom circuit output, and cell voltage were monitored by a Tektronix TDS 9014C oscilloscope. This electronic arrangement is shown in Figure 7.2.⁴

Electrolyte Solutions Solutions were prepared with KCl ⁵ added to distilled, deionized water with typical initial conductivities $\sim 1.2\text{-}1.7 \mu\text{S}/\text{cm}$. Conductivity was measured with an Orion Versa Star meter and Orion SX1-19751 conductivity probe. These solutions were pumped into the CDI cell via an OmegaFlex FPU-423 peristaltic pump. The prototype CDI cell is detailed in Chapter 6.

Procedures

Active-Charging Phase Procedures For initial shielding measurements (ACP),⁶ pulse frequency was 100 Hz, at 50% duty cycle (5 ms pulses), with peak voltage of 1.0 V and

⁴ This electronic arrangement was recorded 7 Aug 2016 (page 130 in the labbook).

⁵ Fisher Chemical P217-500; KCl was chosen for its nearly symmetric diffusivity. While minor differences have been observed with different salts [47], we do not expect salt-specific phenomena.

⁶ See labbook pages 109-116.




Figure 7.2: Experimental setup for observing the collapse of voltage following a rapid-charging period and isolation from the circuit (RP). Green lines indicate wires for the active signal. Orange lines represent (low current) wires used to observe voltages. Red lines are wires power delivery. Black lines are ground. The small boxes represent BNC splitters.

baseline voltage of 0.0 V, as observed on the oscilloscope. Tests were performed with conductivities of 2.5 (2.46-2.54), 35 (35.4-35.7), 510 (509-512), and 3440 (3440-3450) $\mu\text{S}/\text{cm}$. As voltage and current evolved over many orders of magnitude, data was recorded at two cadences (data points taken every 2 μs or 400 ps) for each solution. In each solution-cadence combination, data was taken in (at least) triplicate.

Relaxation Phase Procedures For voltage-collapse measurements,⁷ pulse frequency was 1 Hz, at 0.05% duty cycle (500 μs pulses). Charging voltage was 2.0 V, such that peak-to-peak voltage was 1.0 V, as observed on the oscilloscope which attenuated the signal by 50%. The charging circuit is such that (after an equilibration period) pulses are negative and mean output will be zero. Consequentially, baseline voltage is ~ 0.5 mV and pulse voltage is ~ -0.9995 V.

Tests were performed with conductivities of 37 (37.4-37.5) and 496 (495-497) $\mu\text{S}/\text{cm}$. Data was recorded at four cadences (data points taken every 4 μs , 400 ns, 4 ns, or 400 ps) for each solution. Measurements at the fastest cadences were taken in triplicate.

⁷ See labbook pages 135-137.




Figure 7.3: Fundamental elements of the charging circuit for pulse-charge tests. When $V_{trigger}$ is low, the mosfet is insulative, the capacitor charges from V_{charge} , and V_{out} approaches ground due to the *pull down* resistor. When $V_{trigger}$ is high, the mosfet becomes conductive, pulling the high-side of the capacitor to ground, and the (negative) voltage pulse discharges through anything attached to V_{out} , in parallel with the $100\ \Omega$ resistor. The +15 volt input is required for housekeeping in the mosfet circuit.

7.3 Analysis and Data Processing

Several data-processing steps were required to obtain time evolutions of electric field shielding and internal voltage collapse from our raw measurements of external current and voltages. In this section, we outline these steps, specifically building the datasets, correcting for the Pearson current monitor, applying the charge-current equation, and validating the analysis method. We will illustrate each step through the example of our active-charging phase data for the $2.5\ \mu\text{S}/\text{cm}$ solution.

Building datasets (both ACP and RP)

Building functional datasets required several steps. First, triplicate data was aligned in time. For ACP datasets, we aligned the time of maximum value of the current spike in each set. For RP datasets, we aligned the minimum value on the trigger voltage. With data aligned in time, we then took the mean of each triplicate data point.

Next, data from different cadences were temporally aligned and then merged. The data of interest spans about eight orders of magnitude. This provides higher temporal resolution near the transition of interest (either the initial application of voltage or the effective isolation), and lower resolutions as the curves evolve. This can be seen in our illustrative dataset in

Figure 7.4, where the data is combined from two cadences, each covering \sim four orders of magnitude.




Figure 7.4: Temporally aligned and merged datasets for source voltage, ϕ_{a-e} (blue), and external current, i_e (red), from active-charging phase measurements with $2.5 \mu\text{S}/\text{cm}$ solution. Each curve represents $\sim 20,000$ data points, each the average of points from three pulses. The transition between the two cadences is visible at $4 \mu\text{s}$. We also note the observed voltage rises almost immediately and holds near 1 Volt until $\sim 5 \text{ ns}$; this is roughly the time it takes for the electric field to develop in the ~ 6 feet of BNC cable in the CDI-cell circuit. At that point, the circuit appears as a near-short, suppressing voltage until sufficient charge is transferred.

Accounting for Pearson monitor (ACP only)

The Pearson 6027 is a high-speed current monitor, designed to capture brief transient currents. The monitor's speed is necessary to capture the rapid response, but it has a drawback: the output signal will decay if monitoring a constant or slowly evolving current. It decays via a nearly exponential curve, with a published *droop rate* of $0.2\%/ \mu\text{s}$, [65] equivalent to $\tau_{1/e} = 0.5 \text{ ms}$. Our independent characterization⁸ showed a very good fit to exponential decay, but with $\tau_{1/e} = 1.59\text{-}1.65 \text{ ms}$.

Typically, this decay rate is characterized to place a limit on the accuracy as a function of time (or frequency). Understanding the decay, however, also allows us to correct for it,

⁸See lab book, p.115-117.

and get usable data beyond the ms timescales. The output of the Pearson monitor can be understood as the sum of changes in current, each decaying exponentially, or

$$i_{out}(t_2) = i_{out}(t_1)e^{\left(\frac{t_1-t_2}{\tau_{Pear}}\right)} + \int_{t_1}^{t_2} \frac{di_{in}(t)}{dt} e^{\left(\frac{t-t_2}{\tau_{Pear}}\right)} dt \quad (7.1)$$

where i_{in} is the current through the monitor, i_{out} is the output of the current monitor and τ_{Pear} is the Pearson current monitor decay rate.

By marvelous serendipity, this equation has the same functional form as the charge-current equation (see Equation 5.38). The derivation presented in Chapter 5, therefore, implies a corollary to our original differential equation (5.26). This corollary can be written

$$\frac{di_{in}(t)}{dt} = \frac{di_{out}(t)}{dt} + \frac{i_{out}(t)}{\tau_{Pear}} \quad (7.2)$$

which can then be integrated to go give the more practical

$$i_{in}(t_2) = i_{in}(t_1) + [i_{out}(t_2) - i_{out}(t_1)] + \frac{[q_{out}(t_2) - q_{out}(t_1)]}{\tau_{Pear}} \quad (7.3)$$

In the abstract, this correction equation could be used to correct Pearson output indefinitely. Unfortunately, as i_{out} collapses, the uncertainty due to oscilloscope binning is effectively magnified and eventually dominates the signal. Practically, it extends the reliable data range by a few orders of magnitude. This correction method is shown for our example dataset in Figure 7.5.

Applying the Charge-Current Equation (ACP only)

Resolving timescales for ionic shielding (and double-layer growth) from external current, $i_e(t)$, requires getting $q_c(t)$ (and $q_i(t)$) by applying the charge-current equations, 5.38 and 5.39, both derived in Chapter 5 and reproduced here.

$$q_c(t_2) = q_c(t_1)e^{\left(\frac{t_1-t_2}{\tau}\right)} + \int_{t_1}^{t_2} i_e(t)e^{\left(\frac{t-t_2}{\tau}\right)} dt \quad (7.4)$$

$$\tau \equiv \frac{\epsilon_0 \epsilon_r}{\kappa_b}$$

$$q_i(t_2) = q_i(t_1) + \frac{1}{\tau} \int_{t_1}^{t_2} q_c(t) dt \quad (7.5)$$




Figure 7.5: Correcting the Pearson 6027 current monitor’s signal decay. The lower, tan line shows data prior to correction – the decay at \sim ms is due to the Pearson current monitor. The upper red line shows data after correcting for that decay. As the correction operates on an increasingly small value, the uncertainty rises and we see an erroneous upward trend in the corrected curve. Nonetheless, this correction extends usable data by several orders of magnitude. The lighter red line shows the fit to the corrected data through Equation 7.6. The data shown is from active-charging phase measurements, $2.5 \mu\text{S}/\text{cm}$ solution.

Unfortunately, the integrated exponential form in Equation 7.4 introduces a complexity. When applied to discrete data, the output of such a function is dependent on the cadence of the data. Not only would this create discontinuities at the points where data of different temporal resolution meet, it means that neither regime generates strictly accurate values. To get around this, we can fit our external current data to an analytic function, substitute that function for $i_e(t)$ in Equation 7.4, thus generating analytic functions for $q_c(t)$ and, via Equation 7.5, for $q_i(t)$.

Specifically, our corrected external current data was fitted to a function of the form

$$i_e(t) = A + Be^{\left(\frac{t}{10^{\sigma}}\right)} + Dt \quad (7.6)$$

where A , B , C , and D are fitting parameters. The fit for our corrected data is shown in Figure 7.5. Given a fit of this form, $q_c(t)$ becomes the somewhat daunting

$$q_c(t_2) = \left(q_c(t_1) - A\tau - B\left(\frac{C\tau}{C-\tau}\right)e^{\left(\frac{-t_1}{C}\right)} - D\tau(t_1 - \tau) \right) e^{\left(\frac{t_1-t_2}{\tau}\right)} + \left(A\tau + B\left(\frac{C\tau}{C-\tau}\right)e^{\left(\frac{-t_2}{C}\right)} + D\tau(t_2 - \tau) \right) \quad (7.7)$$

As electric field in the bulk, \vec{E} , is simply proportional to q_c , timescales for shielding of \vec{E} can be measured by observing q_c .

We could use Equation 7.5 to generate an (even more tedious) analytic function for $q_i(t)$. In practice, however, it proved much easier to get at $q_e(t)$ through a discrete approximation of the integral of $i_e(t)$, and use the definition of q_c (Equation 5.11) to find $q_i(t)$.

Verification (ACP only)

As detailed in Chapter 5, measurement of supply voltage, $\phi_{a-e}(t)$, allows independent confirmation of functions for q_c and q_i via Kirchoff's Voltage Law. Given our experimental arrangement, Equation 5.41 must be modified to include the resistance and self-inductance of the coaxial BNC cables, giving a modified verification equation

$$\phi_{a-e}(t) = 2\frac{q_e(t)}{C_2} + \frac{q_c(t)}{C_1} + R_1 i_e(t) + L_1 \frac{di_e}{dt} \quad (7.8)$$

where R_1 is explicitly the resistance of BNC cables and contacts, and L_1 is the BNC cable self-inductance. In our sample case, data from 40 ns to 2 ms could be fit with $C_1=320$ nF, $C_2=580$ μ F, $R_1=0.6$ Ω , and $L_1=24$ μ H.⁹ We note that this fit becomes progressively worse as later timescales are included, suggesting that not only is our model internally consistent at early timescales, but that it becomes less so at later timescales, as expected. This confirmation fit for our sample dataset is shown in Figure 7.6.

It would be easy to infer too much from this confirmation, as there are several caveats. First, the fit was strongly dependent on C_1 , L_1 , and $q_e(0)$, but relatively insensitive to values of R_1 and C_2 . Second, this fit represents the best of a whole class of relatively good fits that can be obtained by adjusting both C_1 and L_1 . As a consequence, the values for these parameters should not be taken as results (with the possible exception of $q_e(0)$). Finally, while internal consistency at early timescales and failed internal consistency at later timescales appears consistent with latter diffusion effects, there is no direct evidence that this is the cause.

⁹ These values are broadly within expected ranges. C_1 was expected in nF- μ F range. C_2 was known to be larger, but less than the full (many F) values reached after full diffusion. R_1 was expected to be between 0 and a few Ohms. Finally, $L_1 \sim \mu$ H for a BNC of these dimensions.




Figure 7.6: Confirming the charge evolution with independent source-voltage measurement. The thick blue line represents $\phi_{a-e}(t)$ as captured via the oscilloscope. The thin orange line represents $\phi_{a-e}(t)$ as calculated from charge evolutions, with C_1 , C_2 , R_1 , $q_e(0)$ and L_1 (self inductance of the cables) as fitting parameters. Datasets are from active-charging phase measurements, $2.5 \mu\text{S}/\text{cm}$ solution.

However, the consistency between modeled and observed voltages does appear to provide a validation of the charge-current equation at early timescales.¹⁰ In future efforts, this validation could be made more robust by reducing the number of fitting parameters. Practically, this would require measuring voltage across the cell, rather than at the signal, reducing the impact of the cables and the need for R_1 and L_1 as fitting parameters.¹¹

7.4 Results and Discussion

In this section, we present the charge evolutions for each of the active-charging phase experiments, define the early shielding timescale and present it as a function of concentration. We then present the cell voltage evolutions for the relaxation phase and identify the timescales for collapse of the bulk voltage and double-layer voltages.

¹⁰ The curve includes two inflection points. As such, it would require 5-variables if fitted via a (4th order) polynomial. Our fit, however, is able to reproduce the curve with functional dependence on only three variables, C_1 , L_1 , and $q_e(0)$.

¹¹ RP measurements were performed in this manner, but without fast-cadence data from the start of the charging pulse, application of the charge-current model proved infeasible.

Time Evolution of Charges in Active-Charging Phase

Figure 7.7 shows the time evolution of charges q_e , q_i , and q_c following an applied voltage. Shielding of the electric field in the bulk can be understood by observing q_c , the electronic charge transferred to the electrodes in excess of ionic charge adsorbed in the double layer.




Figure 7.7: Charge evolutions following an increase in applied voltage from 0 to 1 V for four solution concentrations. 2.5 $\mu\text{S}/\text{cm}$ is nearly pure water, where 3440 $\mu\text{S}/\text{cm}$ is brackish. q_e , in red, shows the electronic charge transferred to the cell. q_i , in green, is the total ionic charge adsorbed in the double layer. q_c , in blue represents the difference between the two: the electronic charge in the electrodes unmatched by ionic charge in the double layer. Electric field in the bulk is proportional to q_c .

There are two parameters of use here. First, we can define an *early shielding timescale*, $\tau_{50\%}$, as the time when $q_i = q_c$. This is, equivalently, the time at which the field expected based on $q_e(t)$ would have been reduced by half by screening effects. Second, we are interested in the maximum values of q_c , here referred to as q_c, max . They correspond to maximum electric field in the bulk, and reach a quasi-steady-state value within early timescales. Both parameters can be identified in Figure 7.7, and are shown relative to concentration in Figure 7.8.




Figure 7.8: At left, timescales for early shielding, relative to concentration. The timescale was defined as the time at which $q_i = q_c$, or the time at which the field expected based on external addition of charge has reduced by half. The experimental data (in red) can be fit well ($R^2=0.9995$) to a power law (in blue). The fit exponent of -0.97 suggests a hyperbolic relation, consistent with the theoretical *RC time constant* for R_2C_1 as defined in 5.25. At right, the maximum values of q_c , relative to concentration. The experimental data (red) can be crudely fit ($R^2=0.98$) to a power law (blue).

There is a potential issue of epistemological circularity here. Specifically, the RC time constant, τ , was developed from theory, and applied in interpretation of the experimental data. As defined, our early shielding timescale, $\tau_{50\%}$, is functionally a combination of rise time for i_e and the theory-originated timescale τ . Since i_e rises quickly, $\tau_{50\%}$ is heavily influenced by τ . We are saved, I believe, by the confirmation step: if we had assumed an unrealistic timescale, we would not have found agreement with our independently observed voltage. Therefore, while the similarity between the two should not be taken as revelatory, we nonetheless have some confidence that the values are valid.

Cell Voltage Collapse in the Relaxation Phase

Before examining the relaxation phase data, we can make some theoretical predictions. In pulse-charged CDI, the cell is functionally disconnected after a brief charging phase. In this relaxation phase, $i_e=0$ and q_e is constant. With this constraint, the charge-current equation (5.38) is reduced to

$$q_c(t_2) = q_c(t_1)e^{\left(\frac{t_1-t_2}{\tau}\right)} \quad (7.9)$$

$$\tau \equiv \frac{\epsilon_0\epsilon_r}{\kappa_b}$$

where t_1 is here any time after disconnection. We, therefore, expect the bulk voltage, ϕ_{c-d} , (proportional to q_c via C_1) to decay exponentially with the time constant τ . After several τ , the cell voltage, ϕ_{b-e} , should reduce to only the double-layer voltages $2\phi_{c-b}$.

The experimental results, however, show two voltage collapses, separated in timescale (the voltage time-evolutions for $37 \mu\text{S}/\text{cm}$ and $496 \mu\text{S}/\text{cm}$ are shown in Figure 7.9). Fitting our predicted exponential decay to the experimental data, we find the theory is consistent with the faster of the timescales.

Our first attempt to identify the second collapse involved adding a second exponentially decaying term to the fit. It soon became clear, however, that the second collapse is not exponential in form. We next considered the possibility that the second voltage collapse was due to ionic diffusion into the porous electrode. If true, we can make two predictions. First, the time-evolution of the double layer voltage, ϕ_{e-d} , will have a functional form consistent with diffusive infiltration into the electrode. Specifically,

$$\phi_{c-b}(t) \propto \frac{1}{C_2(t)} \propto \frac{1}{A_t(t)} \propto \frac{1}{d_i(t)} \propto \frac{1}{t^{1/2}} \quad (7.10)$$

where C_2 is the double-layer supercapacitance, A_t is the effective tortuous area acting in the capacitor, and d_i is the electrode depth infiltrated by ions. Based on this, we fit the cell voltage time-evolution to a function of the form

$$\phi_{e-b}(t) = Ae^{(-t/\tau)} + 2B\frac{(0-C)^{1/2}}{(t-C)^{1/2}} \quad (7.11)$$

in which A , B , and C are fitting parameters, where A corresponds to initial bulk voltage $\phi_{d-c}(0)$, B corresponds to initial double-layer voltage ϕ_{b-c} , and C is an equivalent time, capturing the infiltration that occurred before isolation. As shown in Figure 7.9, this functional form is consistent with the experimental data at both concentrations.




Figure 7.9: Collapse of cell voltage, ϕ_{b-e} , from 1V to 0V after effective isolation from the charging circuit. Data is shown for $37 \mu\text{S}/\text{cm}$ (left) and $496 \mu\text{S}/\text{cm}$ (right) concentrations. In each case, red shows experimental data. The green line represents our first fit, based on a single exponential decay of bulk voltage, ϕ_{d-c} . The blue represents our second fit, which incorporates a term likely tied to ionic diffusion into the porous electrode. The strong oscillations at $\sim 10^{-7}$ s are likely due to under-damped conditions due to minimizing R_1 , the large capacitor in our charging circuit, and the self-inductance of our BNC cables.

Second, we can predict that if the faster collapse is the decay of the bulk voltage, then the difference between initial and inter-collapse cell voltages captures the bulk voltage at the time of disconnection. Each such value should be proportional¹² to q_c at the time of disconnection. Therefore,

$$\frac{q_{c,35\mu\text{S}/\text{cm}}(500\mu\text{s})}{\Delta_{\text{fast}}\phi_{e-b,37\mu\text{S}/\text{cm}}} \approx C_1 \approx \frac{q_{c,510\mu\text{S}/\text{cm}}(500\mu\text{s})}{\Delta_{\text{fast}}\phi_{e-b,496\mu\text{S}/\text{cm}}} \quad (7.12)$$

where $\Delta_{\text{fast}}\phi_{e-b}$ is the change in cell voltage over the early collapse. In theory, we can get the q_c values from the active-charging phase measurements at $500 \mu\text{s}$, the charging time in the relaxation phase experiments.

Unfortunately, this expected result does not bear out: the estimates for C_1 from the two comparisons vary by roughly an order of magnitude. Why? There are a number of sources of

¹² Since capacitance varies with concentration only via the relative permittivity of the solution, which is a very weak function of concentration changes at this scale. Therefore, C_1 will be nearly the same for these concentrations.

small variations that have not been incorporated,¹³ but we suspect none of these are sufficient to explain the discrepancy. It seems more likely that the issue arises from the use of different charging electronics in the two arrangements. Implicit in our (second) prediction is the assumption that 500 μs of charging is equivalent in the two experimental arrangements – this is unlikely. In future work, this question could be resolved by a mixed-phase experiment, recording high temporal resolution data around both the beginning of the pulse, as in the ACP experiments, and after the effective isolation, as in the RP experiments.

Conclusions

The experimental work of this chapter, interpreted via theory developed in Chapter 5, has helped illuminate early timescales, both for shielding during the active-charging phase and for voltage collapse following in the relaxation phase. Through the confirmation method, this data also lends strength to the charge-current relation for early timescales.

Specifically, the ACP experiments provide an early shielding timescale, $\tau_{50\%}$, which appears to be inversely proportional to ionic conductivity, κ , $\sim 1.5\tau$, the "RC" time constant for the electrolyte defined in Equation 5.38.

The RP experiments show that voltage collapse occurs via two steps, largely separated in time. The faster collapse is tentatively consistent with bulk voltage collapsing at the "RC" timescale, τ . The latter appears to follow a $1/t^{1/2}$ form consistent with double-layer voltage decreasing due to diffusion into the porous electrode (see Figure 7.9). It has an effective timescale of $\sim 0.3\text{ms}$, which does not appear to be concentration dependent.

Two limitations preclude resolution of timescales larger than $\sim\text{ms}$ with this approach. First, as the current-charge relation neglects diffusion, it cannot be assumed to be valid beyond this range. Second, the signal from the Pearson current monitor decays at that order, introducing increasingly large uncertainties in ACP measurements. Therefore, we can only take from the quasi-steady-state q_c values that that complete shielding occurs on timescales longer than ms. As any ion removal requires an unshielded field, this is consistent with the proof-of-concept experiments where removal continued for 5+ minutes per cycle (see Figure 6.8).

A Case for Simulation

While the relatively simple charge-current equation has provided insight regarding early timescales, it neglects ionic diffusion into the electrode and into the bulk. It also treats the

¹³ Small effects include concentration and temperature dependence of ϵ_r , differences in concentration between ACP and RP experiments, evolution of effective d_z for C_1 .

bulk concentration as constant and uniform, which are known to be problematic assumptions as ion distributions evolve.

Accurately capturing these phenomena in an analytic frame is a daunting task. Similarly, resolving electric fields and ion distributions over small length-scales is experimentally challenging. Instead, we believe such effects might be better captured via numerical simulation. This work is presented in Chapter 8.

Chapter 8

Numerical Simulation for PC-CDI Timescales

8.1 Introduction

In Chapter 4, we presented the case for using capacitive deionization (CDI) to treat brackish water in the developing world. We reviewed the history of CDI and made a case for investigating possible improvements, introducing pulse-charged capacitive deionization (PC-CDI) and optimal power profiles.

In the PC-CDI arrangement, the electrodes are charged by voltage pulses. The pulses are greater than redox potentials, but short enough that the reactions can be largely avoided. The rising edge of the pulse is driven by an outside voltage, similar to traditional CDI; we call this the *active charging phase*. The falling edge of the pulse is driven by internal equilibration within the cell; we call this the *relaxation phase*. The viability of PC-CDI depends on several timescales, both in the active charging and relaxation phases.

Alternately, in the optimal power profile arrangement, CDI is driven via a slowly-evolving cell voltage such that double-layer voltages specifically remain just under the redox potentials. In this arrangement, cell voltage is kept at roughly the sum of redox potential and any voltage that exists across the bulk.

Following the introduction of these arrangements in Section 4.6, the subsequent chapters have pursued the PC-CDI timescales. In Chapter 5, we developed a theoretical model to predict ion motion in the cell and help interpret experimental results. In Chapter 6, we detailed the development a PC-CDI prototype and experimental testbed. In Chapter 7, we presented experimental results, making the case for parallel numerical simulation.

In this chapter, we present our numerical simulations for CDI. Specifically, we begin by outlining the goals and approach for this work. We then detail the model we are using, presenting the spatial discretization, differential equations, and temporal discretization for the planar electrode case. We then describe the changes made to model a system with porous electrodes. Next, we present simulation results, including model validations and new findings. Finally, we present our conclusions about our model, physical phenomena, and the implications for improving capacitive deionization.

8.2 Goals & Approach

In this section, we describe the aims of this effort, our approach, and make comparisons to previous efforts in literature.

Goals

We sought to simulate the evolution of ion concentrations within a CDI cell, both during the active charging phase and the following relaxation phase. Our aims were to validate our model, confirm our theoretical (Chapter 5) and experimental (Chapter 7) work, and then to resolve and understand timescales not discernible via those methods. Specifically, we aimed to confirm values for $\tau_{50\%}$, the timescale for early shielding defined in Section 7.4. We also sought to replicate and better understand the two voltage-collapse timescales observed in experimental measurements of the relaxation phase.

Approach

Broadly, we have taken the Poisson equation (Equation 5.4) and Nernst-Planck equations (Equation 5.5) as our governing equations, obtaining solutions at discrete points in space and stepping through time, using approximations (of varying orders) to these governing equations. In this respect, it is similar to work published by Morrow, et al. [91]. Unlike Morrow, we have simulated CDI cells with both planar and porous electrodes. The planar-electrode simulations provide a direct comparison to Morrow's results. The porous-electrode simulations extend these results to usable comparisons with experimental work and practical results. Comparing the two offers validation of our assumption of negligible diffusion into the electrode at early timescales.

Our porous electrode simulations have more in common with work by Bazant, et al., who simulated a porous electrode via a constant diffusive flux at the simulation boundary [12]. Unlike Bazant, however, we model the evolving ion concentrations throughout the electrode, providing rigorously motivated and time-evolving ion flux values.

Initial conditions are defined via ion concentrations. Our simulations consider binary electrolytes.¹ Initially, molar cation and anion concentrations are uniform and equal, at $c(0)$ over the bulk. For planar electrodes, they are 0 in the ion exclusion zones. For porous electrodes, they are $nc(0)$ in the electrodes, where n is *porosity*.² As a consequence, initial ion-generated electric fields are 0 everywhere. External electric fields are also 0.

Boundary conditions are dependent on the simulation run. In all simulations, flux across the outer boundaries is zero, though outer boundaries are defined differently for planar- and porous-electrode simulations. Additionally, one of two electrical boundary conditions applies. In the relaxation phase, the cell is disconnected from any charging circuits and the electronic charge on the electrodes is constant; in such simulations, the external electric field is constant (and cell voltage, ϕ_{cell} , varies). In the active charging phase, charging circuits keep the cell voltage constant; in such simulations, the external electric field is varied to keep ϕ_{cell} constant. This is consistent with additional electronic charge being added to the electrode from the charging circuits.

8.3 Planar Electrode Model

We begin with a description of the simulation for the (comparatively simple) planar electrodes. In this section, we detail the division of the planar-electrode CDI cell into discrete slices, the discretization of differential equations, and the method for time-stepping. The complexities associated with porous electrodes will be introduced in the next section.

Modeling the Space

Our CDI cell assumes a small size in one dimension (z), with substantially larger orthogonal dimensions (x, y). We will model this space as continuous and uniform in both x and y , but finite (and non-uniform) in z , effectively approximating the three-dimensional problem with a one-dimensional one.

To solve this problem numerically, we discretize the z dimension, dividing the gap between electrodes into six regions of slices, symmetric about the center plane of the cell (see Figure 8.1). The regions nearest the electrodes contain no ions, and are treated as a single slice.³ In the next regions, a certain distance further from the electrodes, space is discretized into slices of uniform width. Farther out, in the interior regions, slice widths increase exponentially

¹ As it stands, our simulation software can treat any binary electrolyte. Generally, we have modeled KCl, but in comparisons with Morrow, we used an artificially symmetric electrolyte.

² Porosity is commonly represented with ϕ . In this work, I have chosen n to represent porosity to avoid confusion with commonly referenced electrical potentials, which I am denoting with ϕ .

³ When fields are present, there will be a voltage across this ion exclusion zone. This can be seen as an implicit but spatially-motivated Stern voltage.

as they approach the center of the cell. The slice widths are controlled by five specifiable parameters:

- d_z is the width of the entire cell.
- A is here the width of the uniform slices near the electrodes.
- B is here the width of the ion exclusion zones.
- l is here the number of slices in each of the two uniform regions.
- m is here the number of slices in each of the two exponential regions.

The widths of slices in the exponential region are characterized by a factor we here call r , where r can be found by (numerically) solving

$$d_z = 2 \left(B + Al + A \sum_{i=0}^{m-1} r^i \right) = 2 \left(B + A \left(l + \frac{1 - r^m}{1 - r} \right) \right). \quad (8.1)$$




Figure 8.1: Image showing one half of the planar electrode cell volume, divided into discrete slices. The other half is a mirror image about the center line. Each half consists of three distinct regions. Ions are excluded from the region closest to the electrodes (blue), typically with a width, $B \sim 0.1$ nm. The next region is divided into l slices of uniform thickness, A . Typically, $l \sim 100$ and $A \sim 0.1$ nm. The central region includes m slices. Their widths expand via a geometric series to fill the remaining space. m is typically ~ 150 .

The slices are numbered by an index, k , where ion concentrations and charge densities are calculated at integer values of k ($1, 2, \dots$), representing mean values within that slice. Fluxes are calculated at slice boundaries, represented by half-integer indices ($1\frac{1}{2}, 2\frac{1}{2}, \dots$).

For much of this chapter, we will be discussing both continuous and discrete functions. Continuous functions will be specified with parentheses (e.g., $\phi(z)$), while discrete functions will be specified with subscripts (e.g., ϕ_k). In general, the index i will denote ion-specific quantities, where j and k will represent spatial indices.

Differential Equations

In the bulk electrolyte, the evolution of this system will be governed by two equations. Electric field is, with appropriate boundary conditions, given by Poisson's equation,

$$-\vec{\nabla}^2 \phi = \vec{\nabla} \cdot \vec{E} = \frac{\rho_q}{\epsilon_0 \epsilon_r} \quad (8.2)$$

where ϕ is electric potential, \vec{E} is electric field, ρ_q is net charge density, ϵ_0 is the *permittivity of free space* (or the *electric constant*), and ϵ_r is the *relative permittivity* of the electrolyte. Expanding the gradient and simplifying for our one-dimensional space, Poisson's equation becomes

$$-\frac{\partial^2}{\partial z^2} \phi = \frac{\partial}{\partial z} E_z = \frac{\rho_q}{\epsilon_0 \epsilon_r} \quad (8.3)$$

where E_z is the component of \vec{E} in the z axis.

Coupled with this, the time-evolution of ion concentrations are given by Nernst-Planck equations

$$\frac{\partial c_i}{\partial t} = \vec{\nabla} \cdot \left[D_i \vec{\nabla} c_i - c_i \vec{u} + D_i c_i \frac{z_i e (\vec{\nabla} \phi + \frac{\partial \vec{A}}{\partial t})}{k_B T} \right] \quad (8.4)$$

evaluated for each ion species i , where c_i is the concentration, D_i is the *diffusivity*, and z_i is the charge of i ; \vec{u} is the fluid velocity, e is the (positive) *elementary charge*, k_B is the *Boltzmann constant*, T is temperature, and \vec{A} is the (magnetic) vector potential. The term in brackets is total ion flux, and is due to three terms: diffusion, fluid flow, and *electrophoresis* (movement of ions due to an electric field), respectively. As we will treat the fluid as static, and magnetic fields as minimal (no outside magnetic fields; fields due to ion currents negligible), Nernst-Planck equations for our one-dimensional case can be simplified to:

$$\frac{\partial c_i}{\partial t} = D_i \frac{\partial^2 c_i}{\partial z^2} + \frac{D_i z_i e}{k_B T} \frac{\partial}{\partial z} \left(c_i \frac{\partial \phi}{\partial z} \right) \quad (8.5)$$

Applying first-order approximations of the derivatives to this simplified version, gives discrete Nernst-Planck equations

$$\frac{\partial c_{i,k}}{\partial t} \cong D_i \left[\frac{\left(\frac{c_{i,k+1} - c_{i,k}}{z_{k+1} - z_k} \right) - \left(\frac{c_{i,k} - c_{i,k-1}}{z_k - z_{k-1}} \right)}{z_{k+1/2} - z_{k-1/2}} \right] + \frac{D_i z_i e}{k_B T} \left[\frac{c_{i,k+1/2} E_{z,k+1/2} - c_{i,k-1/2} E_{z,k-1/2}}{z_{k+1/2} - z_{k-1/2}} \right] \quad (8.6)$$

where terms represent net diffusive and electrophoretic fluxes across the slice boundaries, respectively.

Concentrations at half-integer indices are approximated by a weighted average

$$c_{i,k+1/2} \cong \frac{(z_{k+1/2} - z_k)c_{i,k+1} + (z_{k+1} - z_{k+1/2})c_{i,k}}{z_{k+1} - z_k} \quad (8.7)$$

which collapses to a simple average in the uniform region.

For the uniform region near the electrodes, with evenly spaced slices, this constitutes a central approximation. As such the first order approximations are good to second order (have only third order corrections) in the crucial near-electrode space.⁴

Electric fields can be found by integrating Poisson's equation (from either electrode), effectively adding the planes of charge.

$$E_z(z_2) - E_z(z_1) = \int_{z_1}^{z_2} \frac{\sum_i c_i z_i}{\epsilon_0 \epsilon_r} dz \quad (8.8)$$

In discretized space, this becomes a sum (again from either electrode),

$$E_{z,k+1/2} \cong E_{z,0-1/2} + \sum_{j=0}^k \frac{(\sum_i c_i z_i)(z_{j+1/2} - z_{j-1/2})}{\epsilon_0 \epsilon_r} \quad (8.9)$$

where $E_{z,0-1/2}$ is any electric field external to the bulk, due to electronic charge on the electrodes. This uniform (due to our geometry) field is added to the non-uniform variable field due to mobile ions. This field is constant in simulations where the cell is isolated from charging circuits (relaxation phase), and non-constant in simulations where charging circuitry keeps cell voltage constant (active charging phase).

Cell voltage can be found by summing (both internal and external) electric fields over the electrode gap.

$$\phi_{cell} \cong d_z E_{z,0-1/2} + \sum_{k=0}^{2(n+m)-1} (z_{k+1/2} - z_{k-1/2}) E_{z,k} \quad (8.10)$$

Time Steps For each forward time step, the change in concentration of each ion, i , for each slice, k , is integrated via

$$c_{i,k}(t + \Delta t) \cong c_{i,k}(t) + \int_t^{t+\Delta t} \left(\frac{\partial c_{i,k}}{\partial t^*} \right) dt^* \quad (8.11)$$

where the integrand is made explicit in Equation 8.6, and within that function, concentrations at half-integer indices are obtained from Equation 8.7 and electric fields are obtained from Equation 8.9.

⁴ There is a tradeoff, as higher-order approximations would lower diffusive error, but may introduce instabilities. As such, the simulation code was written with the option of using higher-order approximations, but we saw very little difference in results, and the option was eventually removed in later versions.

After several such steps, Equation 8.11 will have generated a time-series for concentration of each ion, i , in each slice, k . The most recent values in that time-series could be used to make a polynomial approximation (of arbitrary order, given sufficient history), which could then be used to approximate the forward value. Initially, we used such polynomial extrapolations to generate our first approximations, then iterated until the solutions converged. Eventually, however, we realized that simply beginning with the values at time t (equivalent to a 0th order polynomial approximation) was sufficient for small Δt , and ran faster than higher-order approximations with longer timesteps.

8.4 Integrating Porous Electrodes

While the planar electrode simulations were useful for validating our basic model and building simple understanding, meaningful results require including porous electrodes. In this section, we present our conceptual approach for integrating these electrodes, detail our mathematical treatment, and discuss effects of pore-size distributions and morphology.

Approach and Conceptual Treatment of Porous Electrodes

I first note that despite the complexity of the electrode structure, it is likely the model can still be limited to one spatial dimension, z , allowing the electrode space to be divided into slices comparable to the bulk. With pores at $\leq 10\mu m$ scales and orthogonal electrode dimensions at $\sim mm$ scales, careful treatment can eliminate elements in x and y dimensions.

To accomplish this, we will rely on the assumptions that the electrode material is both isotropic and homogeneous at scales \gg pore sizes. Using these assumptions and symmetry arguments, we first consider which ion fluxes can be reasonably assumed to have zero net effect. Note that although the porous spaces almost certainly have complex geometries, we will still refer to them as *pores* and refer to the average distance from the pore center to the electrode as a *pore radius*.

First, isotropy allows us to assume an unbiased distribution of pore-wall orientations, within the entire electrode and within each slice. Therefore, any effect that generates ion flux toward (or away from) pore walls will have no net effect on movement in z . In contrast, effects generating flux parallel to pore walls may not be cancelled by an equal and opposite flux. Second, ion fluxes parallel to pore walls will have non- z components. For an isotropic materials, these too will cancel. Therefore, only the z -components of parallel-to-pore-wall fluxes will contribute to net ion movement.

As the model tracks concentration as a single value (per ion species) for each slice, this must represent a mean concentration over the slice. Due to local variations in the porous path, there will certainly be small areas where the diffusion front has progressed more or

less. That said, we expect those differences to be small (given the difference in electrode and pore scales) and self-limiting (as substantive non- z gradients will tend to resolve via non- z diffusion). The consequence is that we treat the diffusion-inducing derivative in z as largely superimposed on the porous structure. This derivative too is then a mean value over the slice.

As we will see, it is conceptually useful to consider the z -dimension as three distinct regimes. In the electrolyte, electrophoresis and diffusion are both important and proceed via the treatment presented in the previous section. In the electrode, however, we take the electronic conductivity of the electrode to be substantively greater than the ionic conductivity of the electrolyte. Therefore, deep within the electrode,⁵ we expect negligible electric field effects and ions will move via diffusion alone, slowed by the porous structure. However, in the transition regime near the surface, attenuated electric fields will be present, and limited double-layers may form at pore walls. Here, ions may move by (reduced) electrophoresis in the core of the pore, as well as diffusion both through the core of the pore or parallel to surface walls within the double layers.⁶ These conceptual regimes are illustrated in Figure 8.2. While these three regimes are conceptually useful, we will mathematically describe the deep electrode and transition via continuous functions modifying the relevant terms.

Mathematical Treatment

If we are to compare with experimental results, we will also need to define terms of interest for our model in terms of measurable electrode properties. To begin, we define A_{scs} , the *specific cross-sectional area*, as the area of the intersection of pore volumes with a plane, per area of that plane. This factor then modifies the area available to diffusive flux through that plane, and ranges from 0 to 1. Perhaps unintuitively, for isotropic and homogeneous geometries, A_{scs} is equal to porosity, n , for all planes.

Deep Electrode Diffusion In the deep electrode, ions are concentrated in, and equally distributed throughout, the pores. The local concentration for an ion species in the pore, here called *effective concentration* $c_{i,e}$, is then

$$c_{i,e} = \frac{c_i}{n} \quad (8.12)$$

The derivative of this effective concentration will govern diffusive flux, but the flux will flow through a surface area reduced by the same factor, and the effects will cancel.

⁵ This can be compared to the concept of *optically tight* shielding used for EMF-sensitive electronics. Typically, openings are considered tight when the length \gg diameter. In our case, we can expect shielding to be effective when depth \gg pore diameter

⁶ Note that electrophoresis parallel to surface walls immediately adjacent to the walls must be taken to be zero as electric fields will be normal to the surfaces of a conductive electrode.




Figure 8.2: Cross-sections of the porous electrode (tan), illustrating our three regimes. In the *deep electrode* (left), electric fields are fully shielded by electronic charge in the conductive electrode, and ions move through diffusion only in the porous fraction of electrode. In the *transition*, partially attenuated electric fields induce limited electrophoresis in pore centers. In the *electrolyte* (right), ions move via diffusion and electrophoresis within an electric field which is largely uniform in x and y .

The functional direction of the flux will also be relevant. We have already argued that we should consider only ion flux parallel to pore surfaces. For a (locally planar) surface not aligned with the z -axis, there are two effects. First, the concentration gradient in z will be spread over a longer distance, slowing diffusion. Second, the diffusive flux parallel to that surface will have a smaller component in the z direction. Formally, if we define θ as the angle between the z -axis and a the normal vector of the surface, then each effect introduces a factor of $\sin \theta$. In an isotropic material, the net effect is simply the mean effect over all possible surface area orientations. Effective diffusion, $D_{e,i}$ is then

$$D_{e,i} = D_i \frac{\int_0^\pi \int_0^{2\pi} \sin^3 \theta d\psi d\theta}{\int_0^\pi \int_0^{2\pi} \sin \theta d\psi d\theta} = D_i \frac{2}{3} \quad (8.13)$$

where ψ taken over 2π and θ taken over π represent all angles, and $\sin \theta d\psi d\theta$ represents the solid angle differential in spherical coordinates.⁷

Implicit in this treatment is an assumption about the morphology of the material. Independent of porosity, the "pores" may resemble tubular tunnels through a solid bulk, or they

⁷ The solid angle differential is more frequently represented as $d\Omega = \sin \theta d\phi d\theta$, but within this work, we have reserved the use of ϕ for electric potential.

may be the spaces between similarly tubular filaments. We can informally refer to the former case as *concave* and the latter case as *convex*. In the concave limit, only a single vector is parallel to all adjacent surfaces. In this case, the flux is more constrained, and the corollary integral introduces a factor of $\frac{1}{3}$, instead of $\frac{2}{3}$. Similarly, in convex cases, the equivalent factor rises above $\frac{2}{3}$. Our treatment above assumes the pores are neither convex nor concave. This would correspond to parallel planes of electrodes, or to a more complex shape with an equal fraction of convex and concave elements. For comparison to experimental work, the form should be chosen based on the structure of the electrode being modeled. Based on SEM images, we believe our prototype electrode material may be slightly concave, but without a basis for measuring this directly, we have elected to use neutral concavity value.⁸

Diffusion in the Transition In the transition, weak double-layers have formed, concentrating counter-ions at the surfaces. These ions can move through the pore through electrophoresis or diffusion, or diffuse along the surfaces. Thankfully, net diffusion parallel to pore walls is not affected by the orthogonal concentration gradients. Hence, diffusion in the transition can be treated like diffusion in the deep electrode.

Electrophoresis in the Transition In the deep electrode, pores are at a uniform electric potential. In the transition, however, there exists a voltage drop between electrode and pore centers. This allows for weak double layers. It also allows for voltage gradients in the core of the pore, parallel to pore walls, which will induce electrophoresis. We approximate this potential as

$$\phi_{p-e}(z) = [2\phi(r_p) - \phi_c - \phi_e]e^{\left(\frac{z}{2r_p}\right)} \quad (8.14)$$

where $\phi_{p-e}(z)$ is the voltage between the electrode and the center of the pore at depth $-z$ into the electrode, r_p is the pore radius, $\phi(r_p)$ is the potential at $z=r_p$ into the electrolyte, ϕ_c is the potential at the cell center, and ϕ_e is the potential of the electrode.⁹

Practically, this must be computed as a weighted average based on the electrode's pore-size distribution. For our simulations, we took pore distributions from Gabelich [47], as we are using electrodes with similar material properties from the same manufacturer.¹⁰

⁸ Though we have not rigorously examined the possible connection, this concavity-based factor may correspond to a mean tortuosity, sometimes invoked in describing diffusion through porous media.

⁹ This approximation is based on several concepts borrowed from shielding design for RF-sensitive electronics. The exponential form derives from the fractional attenuation provided by each thin layer of material. As this attenuation is dependent on pore size, $2r_p$ was taken as the characteristic length. In reality, attenuation through long tubes is also dependent on frequency, but that has been neglected here as electric fields become quasi-steady in short order. The somewhat unintuitive coefficient in brackets gives the potential drop between the pore center and electrode at the opening of the pore, and results from both a) limiting gradients into the pore to no greater than those observed in the bulk, and b) requiring a continuous second derivative of ϕ with respect to distance from the pore center, r , at the pore center.

¹⁰ This correction was introduced only in later versions of our simulation code. As such, some of the results shown in the following section do not account for this effect.

Additionally, this electrophoretic flux should correspond only to ions outside the double-layers. This consideration requires estimating non-double-layer cross-sectional area and pore concentrations over a distribution of pore radii and varying Debye lengths. Moreover, the calculations again invoke concavity considerations for which we have no robust measure. Though we have done considerable work exploring this consideration, it has not been implemented in our models to date.

8.5 Results and Discussion

In this section, we present selected results from numerical simulations of CDI cells. Specifically, we show our close replication of planar-electrode results published by Morrow, and outline concentration-shielding time and gap distance-shielding time relationships observed in the planar-electrode simulations. We then compare planar- and porous- electrode simulations with experimental results in the active charging phase. Finally, we use the porous-electrode model to simulate the active charging and relaxation phases.

Validating the Planar Electrode Version

Before integrating porous electrodes, we sought several confirmations that our comparatively simple planar-electrode simulation could produce reasonable results. First, Gouy-Chapman Theory provides steady-state solutions for the electric field,¹¹ given planar electrodes (see Equations 5.7, 5.8, and 5.9). All constant-voltage planar-electrode simulations should, therefore, converge on the corresponding Gouy-Chapman solutions. To make these comparisons easier, our software generated the steady-state solution for each simulation. We have yet to see an relevant evolution that does not tend toward the corresponding solution.

We next sought to replicate the results of Morrow, et al., who simulated planar electrodes by a similar method [91]. Specifically, they simulated a cell with $d_z=1\mu\text{m}$, with a cell voltage brought rapidly to 0.1V for 10 μs and then to 0.0V (shorted) for another 10 μs . We chose this particular effort as it was one of the only ones that simulated a full cell with both electrodes. Figure 8.3 shows the close correlation between their simulation and ours, for the time evolution of electric field in both the bulk and the double-layer. In many ways, duplicating their work here served as a jumping-off point for our simulation work.

Linear Relationships Observed in Planar-Electrode Simulations

As simulations with low concentrations, low cell voltages, or large electrode gaps evolve slowly, we investigated trends with respect to these factors. If basic and robust relationships could be determined, it could justify interpolation or even extrapolation, allowing for faster

¹¹Given certain assumptions, this also effectively provides a profile for ion concentrations in simple electrolytes.




Figure 8.3: Comparing our simulated time-evolution for electric field with those published in the Journal of Physics [91]. The solid black line shows Morrow’s double-layer field; the red line shows our equivalent values. The dashed black line shows Morrow’s central (bulk) electric field; the blue line shows our equivalent values. The small deviations are likely due to different functional forms for voltage rise and different slice widths. In both cases, the simulated cell has planar electrodes with $d_z=1 \mu\text{m}$, charged to cell voltage of 0.1V for 10 μs , then brought to 0.0V (shorted) for 10 μs . $T = 298 \text{ K}$, $\epsilon_r = 78.5$, and both ions had the (relatively artificial) diffusivity, $D_i = 2.03 \times 10^{-9} \text{ m}^2/\text{s}$. There was no ion exclusion zone ($B=0$), and the width of the first slice (A) was $\sim 0.3 \text{ nm}$.

equivalent simulations. We began looking at timescale for final shielding of the electric field. We defined this timescale, τ_s , as the time at which the electric field in the center of the cell has been reduced from its maximum value by a factor of e , or

$$E_{z,c}(\tau_s) = \frac{E_{z,c,max}}{e} \quad (8.15)$$

When electrolyte concentration and applied voltage are held constant, we observe a linear (nearly proportional) relationship between electrode gap, d_z , and shielding timescale, τ_s . Further, the time-evolutions for the central electric field appear to vary only by a scaling factor in time – the functional form is self consistent. These effects are illustrated in Figure 8.4, over two orders of magnitude in τ_s .




Figure 8.4: Varying electrode gap, d_z . At left, decay of the central electric field (due to ion shielding) under a constant cell voltage. From left to right, the d_z was 10 μm (red), 30 μm (orange), 100 μm (green), 300 μm (blue), and 1 mm (violet). In all cases, electrolyte concentration was 2 g/L KCl and applied cell voltage was 1 V. At right, data for shielding timescale, τ_s as a function of electrode gap, d_z . Simulation data (blue points) are well-explained ($R^2 = 0.999$) by a linear fit (green line) with time-axis intercept of 90 ns, or approximately the time of maximum central electric field.

When electrode gap and applied voltage are held constant, we observe an additional, inverse relationship between electrolyte concentration, c_{KCl} , and shielding timescale, τ_s .¹² Again, we see a self-consistent functional form in the decay. As with changes in electrode gap, changes in electrolyte concentration appear to vary the time-evolutions for the central electric field only by a scaling factor in time. The time-evolutions and timescale trend are illustrated in Figure 8.5, again spanning two orders of magnitude in τ_s .

Capturing these relationships provided significant flexibility in designing simulations. Specifically, conditions with equal values of the quantity ($d_z c_{KCl}$) lead to equivalent results. This allows us to approximate slower-running simulations (high d_z , low c_{KCl}) with more rapidly running ones (low d_z , high c_{KCl}), allowing us to decrease effective processing time for some simulations. While useful, these relations are limited to the planar electrode, constant voltage case. As we will see, the inclusion of porous electrodes introduces additional complexities, and we have not yet discovered equivalent relationships in the porous-electrode simulations.

¹² When running these simulations, we used the previously discovered relation between electrode gap and shielding timescale to adjust data from simulations with varying gap sizes to a consistent equivalent electrode gap.




Figure 8.5: Varying electrolyte concentration, c_{KCl} . At left, decay of the central electric field (due to ion shielding) under a constant cell voltage. From left to right, c_{KCl} was 50 g/L (red), 5 g/L (orange), 2 g/L (green), 1 g/L (blue), 500 mg/L (violet), and 50 g/L (gray). In all cases, data has been corrected to an equivalent electrode gap of 1 mm, and applied cell voltage was 1 V. At right, data for shielding timescale, τ_s as a function of electrolyte concentration, c_{KCl} . Simulation data (blue points) were fit ($R^2=0.999$) to a power law with exponent -0.929, suggesting an inverse relationship.

Comparing Planar-Electrode Simulations, Porous-Electrode Simulations, and Experimental Results

In our theoretical formulations (Chapter 5), we had predicted the early phenomena would be dominated by electrophoretic motion, and that diffusion—both into the bulk and back into the electrode—would prove negligible at early timescales. Indeed, this was the justification for neglecting diffusion in our derivation of the Charge-Current equations (see Equations 5.38, 5.39, derived in Section 5.4). If prediction is valid, we expect that planar-electrode simulations, which effectively disallow diffusion into the electrode, should reproduce experimental results at early timescales. Porous-electrode simulations, in contrast, should be able to reproduce experimental results at all timescales.

To test this, we ran planar- and porous-electrode simulations of active charging phase for the $35 \mu\text{S}/\text{cm}$ case presented in Section 7.3. The comparison can be seen in Figure 8.6. As expected, the planar-electrode simulation reproduces experimental data at early timescales, but diverges at late timescales. Unfortunately, the porous-electrode simulations diverge at roughly the same timescales, albeit in a different manner. Further complicating the

interpretation, the divergence occurs at roughly the timescale of the correction made for the Pearson current monitor. Indeed, by adjusting the Pearson timescale used in that correction (within the uncertainty for that value), the experimental data can be made to roughly match either simulation. Thus, we have mixed conclusions: while the early agreement between planar-electrode simulations and experimental data validates neglecting diffusion in early timescales, we cannot resolve whether planar- or porous-electrode simulations provide a better fit to late-timescale experimental data.




Figure 8.6: Numerical simulations of experimental results in the active charging phase. Simulation cell voltage (dark blue) was programmed to match experimental cell voltage (light blue). External current, i_e , was treated as the independent variable. Experimentally measured current (light red) is well replicated by both planar-electrode simulations (green) and porous-electrode simulations (black), until the last decade. At timescales longer than $\sim 10^4$, however, our experimental current is highly dependent on the Pearson correction (see Section 7.3). At left, the experimental current is interpreted with our best experimentally determined value for the Pearson timescale (1.62 ms). At right, the experimental current has been adjusted to fit the porous simulation by modifying the Pearson timescale to 1.24 ms. This difference in Pearson timescales is far smaller than the discrepancy between our experimentally determined value (1.59-1.65 ms) and that specified by the manufacturer (0.50 ms). Unfortunately, this means that our experimental data is likely insufficient to meaningfully endorse either simulation at late timescales.

Time Evolution of the Active Charging Phase

After validation steps, we sought to capture timescales in the active charging phase. In these simulations, cell voltage is always specified, beginning at zero, quickly rising to an asymptote, and then remaining effectively constant. This is comparable to traditional CDI. These simulations allow for determination of $\tau_{50\%}$, and with sufficient run-time, τ_s . Concentration and electric potential data from an illustrative simulation is shown in Figure 8.7.¹³

Of particular note are the flattening of the ion concentration profiles in the electrode following establishment of the double layer. This phenomenon is most visible from 20 μs to 200 μs and from 200 μs to 2ns in the example shown. This provides confirmation of nuanced behavior not captured in the Bazant simulations, which assume a constant diffusive flux into the electrode [12].

Time Evolution of the Relaxation Phase

Using the porous electrode model, we then simulated CDI cells being charged and then electronically isolated, capturing the relaxation phase. Based on our understanding and experimental data (see Figure 7.9), we expected the cell voltage collapse over two distinct timescales. The first collapse should decay exponentially with timescale τ , the “RC-timescale” for the electrolyte, defined in Equation 5.25. This would correspond to the collapse of the voltage gradient across the bulk. For the second collapse, we expected a $1/t^{1/2}$ form,¹⁴ consistent with diffusion into the electrode.

Simulated relaxation phase cell voltage collapses are shown in Figure 8.8.¹⁵ As expected, they show a two-step collapse, and the first collapse is indeed an exponential decay consistent with τ . The second collapse, however, has an atypical form, and is not easily fit to simple functions.

As shown, we varied the length of time spent in active charging prior to relaxation. We see the fraction of the collapse occurring at the first timescale decreases with the amount of time spent charging. This is of critical importance, as $\tau \sim \text{ns}$ or smaller for brackish concentrations, while the diffusive collapse is much slower. It is likely that first collapse is fast enough to avoid the instigation of redox reactions,¹⁶ but that the second is not. The

¹³ While we performed such simulations for practical values, the figure shows data from a cell with an artificially small electrode gap (1 μm instead of 1 mm), as this geometry more clearly illustrates some of the features captured in this class of simulations.

¹⁴ Specifically, see Equation 7.11, which includes our fit to the second collapse in experimental data

¹⁵ Again, while we have simulated CDI cells with practical dimensions, the figure shows the artificially small electrode gap. This was done to more clearly illustrate some features of these simulations within tractable run times.

¹⁶ The timescales for instigation of redox reactions vary with reaction and electrode material. For a given combination, these may be captured by impedance spectroscopy and represented in Bode plots, via




Figure 8.7: Active charging phase concentrations (top row) and electric potentials (bottom row), for the whole electrode gap (left column) and near the negative electrode-electrolyte boundary (right column). In both concentration graphs, cation concentrations are solid, and anion concentrations are dotted. Potentials are relative to the potential at the cell's center. In all charts, data is presented at 1 ps (red), 2 μs (orange), 20 μs (green), 200 μs (blue), and 2 ms (violet). Cell voltage asymptote was 100 mV, electrode gap, $d_z=1 \mu\text{m}$, and concentration of 27.5 mM KCl (2050 ppm TDS). At the upper left, we can observe depletion in the bulk, and the local decreases in concentration first presented by Bazant [12] (compare with Figure 5.4), now integrated with ion densities through the electrode. At upper right, we can see the formation (red-blue) and then stabilization (blue-violet) of the electric double layer. At lower left, we see the decaying slope of ϕ , the driving force for electrons in the bulk. At lower right, we see the growing potential drop across the double layer.




Figure 8.8: Relaxation phase simulations of cell voltage collapse. At left, curves show cell voltage following electrical isolation, after charging for 200 ns (orange), 20 μs (green), 200 μs (blue), and 2 ms (violet). These colors correspond to electrical isolation starting at the similarly colored times in Figure 8.7. In all cases, the first collapse is an exponential decay and occurs at a charge-independent timescale $\sim 2\text{ns}$. The second collapse is slower, has a more-nuanced functional form, and a charge-dependent timescale. At right, curves show electric fields at the center of the bulk, $E_c(t)$ (solid blue), and between the electrode and Stern layer $E_{dl}(t) - E_{dl}(final)$ (dotted), for the blue 200 μs charging case. This shows that the first timescale corresponds to the collapse of voltage gradient in the bulk, while the second corresponds to a decrease in field at the double-layer, consistent with ionic diffusion into the porous electrode. As with Figure 8.7, simulation was for a cell voltage of 100 mV, electrode gap of $d_z=1\ \mu\text{m}$, and concentration of 27.5 mM KCl (2050 ppm TDS).

limits on pulse height will therefore be dependent on charge accumulated in the electrode, in addition to the relevant timescales.

Figure 8.8 also shows the time evolutions of electric fields in the relaxation phase. We see that the field in the center of the cell decays to zero, entirely in the first collapse and that the field between the electrode and Stern layer decays to a nonzero asymptote, entirely in the second collapse. After that point, the combined charge in the electrode and double-layers are equal, and ions no longer move through the bulk electrophoretically. The remaining cell voltage then is due to a net charge in the electrode matched by an equal and opposite net ionic charge in the double layer. This will decrease only as excess ions in the double layers move into the porous electrode, a process limited by diffusion.

the frequency at which the impedance phase-shift approaches zero. Typical timescales range from ms to μs orders. While we had anticipated specific characterizations, the second collapse has proven sufficiently slow to render such work unnecessary.

So why then the atypical functional form in the second collapse? As with the active charging phase, We plotted concentrations and electric potentials over the cell for selected times (see Figure 8.9). As we can see, the second collapse results from several phenomena, including diffusion into the electrode, evolving concentration profiles in the electrode and in the bulk, and eventually (in our example, by $\sim 200 \mu$), diffusion from the electrode back into the bulk. As a result, generating a precise analytical form may require additional analysis, though our basic fit to $1/t^{1/2}$ seemed adequate for our experimental data.

While our relaxation-phase simulations display the same features as their experimental counterparts, we were not able to precisely simulate these experiments, as we did for the active charging phase. While detailed voltage and current data was gathered following the isolation, the evolution of voltage was not fully captured during the active phase preceding the isolation. As such, we cannot precisely replicate the charge transferred to the electrodes prior to the relaxation phase.¹⁷ Therefore, it is unclear if the second-collapse data would diverge, or if the collapse is dominated by $\sim 1/t^{1/2}$ phenomena at those time- and length-scales.

8.6 Conclusions

In this section, we describe our assessment of the current model, present our principle phenomenological findings, and discuss the implications for pulse-charged CDI.

Model Validity

To better understand relevant timescales and underlying phenomena, we developed a numerical simulation to model time-evolution of electric potential and ion concentrations. The planar electrode version reproduced results published by Morrow, who had also used numerical simulation to model the planar electrode case [91]. The porous electrode version reproduced features first demonstrated by Bazant, who treated diffusive flux into the electrode as constant [12].¹⁸

We also compared porous electrode model results with our own experimental work. In active charging simulations, forcing a match to our voltage data allowed us to reproduce our external current measurements from ns to 100 μ s scales. In relaxation phase simulations, we were able to reproduce features from experimental data, but direct comparisons were not possible.

¹⁷ This challenge is further complicated by limitations of the simulation software. In the current version, simulating a 1 mm electrode gap porous-electrode cell for 10 ms would require nearly a week. Unfortunately, the software can support continuous simulation for no more than a few days.

¹⁸ As we created a more nuanced treatment of ions in the electrode, and hence, diffusive gradients, we did not attempt to recreate these results precisely




Figure 8.9: Relaxation phase simulations for concentrations (top row) and electric potential (bottom row), for the electrode gap and electrodes (left column) and near the negative electrode-electrolyte boundary (right column). In both concentration graphs, cation concentrations are solid, and anion concentrations are dotted. Potentials are relative to the potential at the cell's center. Data is presented at 1 ps (violet), 20 ns (blue, potential charts only), 20 μs (green), 200 μs (orange), and 2 ms (red) after isolation. This data follows 200 μs of charging at 100 mV (equivalent to the violet line in Figure 8.8). As before, electrode gap, $d_z=1 \mu\text{m}$, and initial concentration of 27.5 mM KCl (2050 ppm TDS). At the upper left, we can observe diffusion into the electrode, as well as increase of ions in the bulk. At upper right, we can see the shrinking double layer. In the lower graphs, we see the bulk voltage collapses completely by 20 ns (violet to blue), and the double-layer voltage shrinks over the later timescales (blue-red).

Based on these results, we believe our simulations to be valid for most cases explored. There are two areas that may require further comparison. First, the active charging phase results diverged at \sim ms scales. We believe uncertainty in the experimental results to be the most likely reason, but this hypothesis requires verification. Second, a direct comparison of the second collapse could help resolve minor discrepancies in the functional form of that feature.

Time Evolution of Ion Concentrations and Electric Potential

The timescales for initial shielding, $\tau_{50\%}$, are consistent with experimental results, and proportional to τ . In active charging phase, the double-layer is first established, then stabilizes as ion removal continues. Final shielding occurs over breakthrough timescales, as bulk potential gradient gradually decreases.

In the relaxation phase, the cell voltage collapses in two stages. In the first, voltage across the bulk collapses to zero via an exponential decay with timescale τ . In the second, double layer voltages collapse to near-zero asymptotes via a complex functional form. The fraction of voltage lost in the first stage is dependent on the ratio of voltage across the bulk to voltage across the double-layers.

Implications for Pulse Charged-CDI

Early on, we had theorized that a very rapid pulse might charge an electrode before initial shielding, thus generating a transient, but substantially higher electric field. This appears to be infeasible, both because this spike does not appear in porous electrode simulations, even with rapid charging, and because $\tau_{50\%}$ is shorter than charging rates achievable with practical electronics.

For PC-CDI to prove viable, the relaxation phase voltage collapse must occur before redox reactions can be instigated. The first stage of collapse is likely fast enough to achieve this, but the second stage is almost certainly not. Therefore, charging pulses must be designed such that the first collapse is sufficient to reduce cell voltage to sub-redox levels. However, as the second collapse corresponds to the collapse in the double-layer, where comparisons to redox potentials are relevant, this is tantamount to requiring that the double-layer voltage never exceeds redox requirements.

Since we cannot exceed redox potentials in the double-layer, there is no reason to limit the additional bulk potential drop to a fractional duty cycle; any additional voltage we apply to create a voltage gradient across the bulk could instead be applied on a continual basis. With this requirement and the duty cycle taken to 100%, PC-CDI is equivalent to the operational mode we termed *optiamal power profiles* in Section 4.6. In this arrangement, total cell

voltage is maximized to keep double-layer voltages just under redox potentials throughout the evolving removal phase. **In essence, the discovered slow collapse of double-layer voltages specifically, in porous electrode systems, means that pulse-charged CDI can never create higher effective ion flux than could the (now seemingly presciently named) *optimal power profiles* arrangement.**

Chapter 9

Conclusions and Closing Remarks

9.1 Introduction

My PhD work has focused on drinking water challenges in the rural developing world. In Chapter 1, I described my personal goals and the motivations for two distinct, but complimentary research paths. In the first path, we have evaluated drinking water interventions as implemented in real-world conditions, seeking to understand what makes such projects successful and robust. This work is presented in Chapters 2 and 3. In the second path, we worked to improve appropriate drinking water treatment technologies in the lab, with the specific aim of improving per-cell flow rates for capacitive deionization (CDI). This effort is presented in Chapters 4, 5, 6, 7, and 8. In this chapter, we summarize the conclusions of these two efforts and offer retrospective thoughts on the methods used in each.

9.2 Conclusions from the Poshiiir River Watershed Study

The efforts for the Poshiiir River watershed case studies are presented in Chapters 2 and 3. In this section, we summarize the study, state conclusions, and offer thoughts on the methods developed and used here.

Study Conclusions

For this effort, we built comparative case studies of single-village piped-water drinking water projects, as implemented, in the Poshiiir River watershed in rural Maharashtra, India. The study included five successful and five failed projects. As an exploratory effort designed to develop and rate potential hypotheses, we sought to identify in the projects and villages that would affect the project's robustness. Data was gathered from disparate sources and compared by unique triangulation method (Chapter 3). These findings are represented in network-based models for each project and corresponding village.

We propose two tentative theories. First and fundamentally, we posit that disruptions to the network can traverse the network, moving across healthy links through healthy nodes, creating additional impacts at those nodes. Even when initial disruptions are detrimental, the resultant impacts can be either detrimental or beneficial (restorative). Second, we theorize that for a scheme to be robust to detrimental disruptions, the network must have healthy feedback loops, capable of responding restoratively to typical detrimental disruptions.¹

We have also pursued typological theories for scheme failures, in which recurring classes of causal chains are identified as types or common features within networks with similar characteristics. While we have identified several candidate causal chains, primarily within the Ainachiwadi scheme, it appears that additional work will be required to develop such theories with rigor.

Thoughts Regarding the Study Methods

Source Scoring, the nascent triangulation method developed for the study and presented in Chapter 3 appears to have performed well, both in our study and in an artificial test case. It has promising potential applications. It provides a method for reconciling data from multiple sources of unknown validity in the absence of a priori data.² However, despite this promise, further evaluation of the method is certainly required.

When paired with process tracing, Source Scoring may be used to quantitatively compare candidate causal mechanisms, in the absence of large statistical sets. If this promise can be realized, these analysis methods may have a greater impact than anything else presented in this work.

While many features of data collection have proven useful, many of the analysis steps were developed after data collection. As a result, there is room for much optimization.³ Specifically, additional comparison questions could be include to make Source Scoring comparisons stronger and more uniform. Added focus on relationships and chronology would make the network-based process tracing more robust.

¹ Recent informal discussions of these ideas have led me to believe that this hypothesis is consistent with a paradigm applied in industrial engineering and manufacturing, where the stabilities of crucial metrics are analyzed with respect to the degree of redundancy in stabilizing feedback loops. In that sector, feedback loops are typically automated and redundant, and those optimizing seek to decrease defect rates over large numbers of trials. As I am only beginning to familiarize myself with these ideas, I hesitate to draw this connection too firmly. However, the more formal application of these well-developed evaluation methods to single high-failure cases may merit attention.

² This can be seen as an alternative to psychometrics, with different strengths and weaknesses.

³Using our newly developed analysis methods on datasets not built with them in mind leaves me with a strong desire to immediately conduct a second study.

9.3 Conclusions from the Pulse-Charged Capacitive Deionization Work

Our work on capacitive deionization was aimed at improving performance by using novel power arrangements to increase effective ion mobility. In this section, we summarize the requirements for such improvements, describe the contributions of each method used, summarize the results obtained for relevant timescales, and assess the viability of pulse-charged capacitive deionization and related operational modes.

Goals and Requirements

Of particular promise was *pulse-charged capacitive deionization* (PC-CDI), an arrangement in which charge is transferred to the electrodes in discrete pulses. In this arrangement, charging occurs in two alternates phases. In the *active charging phase*, the cell is briefly raised to a voltage above redox potentials by an external source. This brief period resembles the beginning of charging in traditional CDI. In the *relaxation phase*, the cell is effectively isolated from charging circuitry, and internal equilibria drives a collapse of cell voltage to below-redox levels.

Practical viability of this method is dependent on a number of timescales. Of specific interest are the timescales for the onset of ionic shielding, $\tau_{50\%}$, and complete ionic shielding, τ_s , in the active charging phase; and the timescales for electrophoretic voltage collapse, τ , and diffusive voltage collapse (which has a more-nuanced, non-exponential form) in the relaxation phase.

In PC-CDI, the effective pulse duration will be a sum the time spent delivering charge (a function of external circuitry) and the voltage collapse timescale(s). To avoid redox, pulse duration should be also be shorter than redox instigation times. Additionally, if pulse duration is shorter than initial shielding timescales, we might gain nonlinear benefits.

Comparison of Methods

In our efforts to assess the viability of PC-CDI, we pursued the relevant timescales via analytical methods (Chapter 5), laboratory experimentation (Chapter 7), and numerical simulation (Chapter 8). While our theoretical work did not lead to analytic solutions for the entire system, we were able to develop a theory for determining internal variables of interest from experimentally observable ones. This theory proved crucial in design and analysis of our experimental efforts.

Our experimental work, aided by the theoretical formulations, proved useful for capturing early timescales. Specifically, we were able to capture timescales for the onset of ionic

shielding in the active charging phase, and voltage collapse by electrophoresis and diffusion in the relaxation phase. We believe the latter to be a new result. The functional form of curves fitted to our experimental double collapse data suggested that the initial, exponential collapse corresponds to the bulk voltage and that the second collapse corresponds to diffusive collapse of the double layer. Unfortunately, this was only a weak implication; we still sought verification.

Our numerical simulation efforts provided the richest understanding of the CDI cell. We developed a model which simulated the time evolution of ion concentrations and electric potential, both in the bulk electrolyte and in the porous electrode. To our knowledge, fully integrated porous electrodes had not been simulated previously. The simulations confirmed experimental values for early timescales in the active-charging phase. More importantly, they reproduced features in the relaxation phase. The additional information provided by the simulation allowed us to confirm that the initial voltage collapse occurs over the bulk, and the secondary voltage collapse occurs in the double layers. We believe this too represents new information.

Timescales and Viability of Pulse-Charged Capacitive Deionization

The viability of PC-CDI depends on several crucial timescales. Theoretical, experimental, and numerical methods all suggest the early-shielding timescale, $\tau_{50\%}$, is near the “RC” timescale of the electrolyte, τ , defined in Equation 5.25. For applicable brackish electrolytes, this is on order a few ns, out of the reach of practical charging electronics (see Figure 7.7). While this precludes a theorized class of PC-CDI, it was not requisite for all versions of the arrangement.

The relaxation-phase timescales proved of more direct relevance. The double collapse of cell voltage in the relaxation phase was observed in experiments and simulation. In each case, it consists of a rapid (or order τ) exponential decay and slower, non-exponential, second collapse to a near-zero asymptote. As confirmed in the simulation, the second collapse corresponds to the collapse of the double layer voltage. As the double layer voltage is the one relevant for redox reactions, and decays too slowly to avoid redox instigation, we can conclude that we cannot create a pulse that brings double-layer voltage above redox potentials. With that limitation, there is no reason that a bulk voltage, which dictates ion flux through the bulk, should be maintained intermittently. Any PC-CDI arrangement will do less to improve ion motion than would optimal power profiles (below).

Viability of Other Operational Modes

In optimal power profiles, applied cell voltage has a slowly evolving value, designed to keep double-layer voltages just below the values that would instigate redox reactions. At any time, cell voltage would be approximately the redox potential plus the voltage across the bulk. Practically, this means a large early pulse that decreases as ions shield the bulk from electric fields. Conceptually, this is equivalent to PC-CDI with double layer voltage limited to avoid redox, and duty cycle of 100%. This still appears to represent a feasible improvement to CDI performance, and, to our knowledge, has not been discussed in literature.

Early in the conceptualization of PC-CDI, we had considered voltage-pulsed CDI, where the falling edge of the pulse was driven by shorting the cell. This was shown to be less-efficient than PC-CDI for low voltages. For high voltages, where this method might induce the first Wien effect, it was found to have high current requirements. We therefore deemed this operational mode infeasible for either low or high voltages (see Section 4.6).

9.4 Closing Thoughts

I would like to close by sharing a few less formal thoughts. Unlike what has been presented previously, these ideas are not based in rigorous exploration or well-vetted literature. Rather, they are impressions, capturing my thoughts at the time of writing.

From what I have gathered, early drinking water interventions in the developing world typically introduced new technical assets suddenly, and with little or no follow up. These efforts failed at astounding rates. More recently, groups have sought to maintain support structures, following such implementations, but this often proves costly and sudden withdrawal of support appears to have similar results. In recent years, groups have seen relatively good success partnering with locally based organizations, implementing, and gradually handing over the reins.

A casual glance suggests that sudden changes lead to unpredictable – and often detrimental – results. This is consistent with the model we presented for the Poshiiir watershed schemes, where disruptions traverse networks leading to additional impacts, beneficial and detrimental. It also appears consistent with theories recently formalized by Owen Barder [9] and Ben Ramalingham [121], in which development outcomes occur within complex adaptive systems, best modeled as overdriven networks of nonlinear links. In such systems, development becomes an emergent property.

In such paradigms, the approach of intervention, support, and gradual transfer has advantages. While the rapid introduction may send ripples through the network, the support stabilizes it until the network adapts to the introduced changes. However, the question

remains: can it be done more efficiently? Certainly there are success stories of targeted disruptive technologies, both in business and in development. If the networks could be rigorously understood, could disruptive feedback be minimized at the time of implementation? It's an enticing idea, but one that likely requires much greater understanding of the interaction between infrastructure and the communities interacting with them.

Bibliography

- [1] Marc D. Andelman. “Charge Barrier Flow-Through Capacitor”. Patent CA 2444390 (Canada). 2007.
- [2] Marc A. Anderson, Ana L. Cudero, and Jesus Palma. “Capacitive deionization as an electrochemical means of saving energy and delivering clean water. Comparison to present desalination practices: Will it compete?” In: *Electrochimica Acta* 55 (2010), pp. 3845–3856.
- [3] Bill B. Arnold and George W. Murphy. “Studies on electrochemistry of carbon and chemically modified carbon surfaces”. In: *Journal of Physical Chemistry* 65 (1961), pp. 135–138.
- [4] Anand N. Asthana. “Decentralisation and Supply Efficiency: The case of Rural Water Supply in Central India”. In: *Journal of Development Studies* 39.4 (2003), pp. 148–159.
- [5] E. Ayranci and B. E. Conway. “Adsorption and Electrosorption at High-Area Carbon-Felt Electrodes for Waste-Water Purification: Systems Evaluation with Inorganic, S-containing Anions”. In: *Journal of Applied Electrochemistry* 31 (2001), pp. 257–266.
- [6] Erol Ayranci and Brian E. Conway. “Adsorption and Electrosorption of Ethyl Xanthate and Thiocyanate Anions at High-Area Carbon-Cloth Electrodes Studied by in Situ UV Spectroscopy: Development of Procedures for Wastewater Purification”. In: *Analytical Chemistry* 73.6 (2001), pp. 1181–1189.
- [7] Erol Ayranci and Brian E. Conway. “Removal of Phenol, Phenoxide and Chlorophenols from Waste-Waters by Adsorption and Electrosorption at High-Area Carbon Felt Electrodes”. In: *Journal of Electroanalytical Chemistry* 513 (2001), pp. 100–110.
- [8] Allen J. Bard and Larry R. Faulkner. “Techniques Based on Concepts of Impedance”. In: *Electrochemical Methods: Fundamentals and Applications*. 2nd. 2001, pp. 368–416.
- [9] Owen Barder. *The Implications of Complexity for Development*. Kapuściński Development Lecture. URL: <https://www.cgdev.org/media/implications-complexity-development-owen-barder>.
- [10] Nittin Bassi and M. Dinesh Kumar. *NREGA and Rural Water Management in India: Improving the Welfare Effects*. Occasional Paper. Institute for Resource Analysis and Policy, 2010.

- [11] Theodore F. Baumann et al. “High Surface Area Carbon Aerogel Monoliths with Hierarchical Porosity”. In: *Journal of Non-Crystalline Solids* 354 (2008), pp. 3513–3515.
- [12] Martin Z. Bazant, Katsuyo Thornton, and Armand Ajdari. “Diffuse-Charge Dynamics in Electrochemical Systems”. In: *Physical Review*. E 70 (2004), p. 021506.
- [13] “Bell Research Develops Leak Sensor for Tanks”. In: *Los Angeles Times* (1989). URL: articles.latimes.com/1989-08-22/business/fi-992_1_leak-sensor.
- [14] Andrew Bennett. “Process Tracing and Causal Inference”. In: *Rethinking Social Inquiry: Diverse Tools, Shared Standards*. 2010, pp. 207–219.
- [15] P. M. Biesheuvel. “Thermodynamic Cycle Analysis for Capacitive Deionization”. In: *Journal of Colloid and Interface Science* 332 (2009), pp. 258–264.
- [16] P. M. Biesheuvel and M. Z. Bazant. “Nonlinear Dynamics of Capacitive Charging and Desalination by Porous Electrodes”. In: *Physical Review*. E 81 (2010), p. 031502.
- [17] P. M. Biesheuvel, H. V. M. Hamelers, and M. E. Suss. “Theory of Water Desalination by Porous Electrodes with Immobile Chemical Charge”. In: *Colloids and Interface Science Communications* (9 2015), pp. 1–5.
- [18] P. M. Biesheuvel and M. van Soestbergen. “Counterion Volume Effects in Mixed Electrical Double Layers”. In: *Journal of Colloid and Interface Science* 316 (2007), pp. 490–499.
- [19] P. M. Biesheuvel et al. “Attractive Forces in Microporous Carbon Electrodes for Capacitive Deionization”. In: *Journal of Solid State Electrochemistry* 18 (2014), pp. 1365–1376.
- [20] John W. Blair and George W. Murphy. “Electrochemical Demineralization of Water with Porous Electrodes of Large Surface Area”. In: *Saline Water Conversion*. American Chemical Society, 1960, pp. 206–223.
- [21] J. O’M. Bockris, M. A. V. Devanathan, and K. Müller. “On the Structure of Charged Interfaces”. In: *Proceedings of the Royal Society* 274 (1356 1963), pp. 55–79.
- [22] J. O’M. Bockris, A. K. N. Reddy, and M. Gamboa-Aldeco. *Modern Electrochemistry*. 2nd. Kluwer Academic / Plenum Publishers, 2000.
- [23] William L. Bourcer et al. “Deionization and Desalination using Electrostatic Ion Pumping”. Patent 7981268. 2011.
- [24] Romain Broséus et al. “Removal of Total Dissolved Solids, Nitrates and Ammonium Ions from Drinking Water Using Charge-Barrier Capacitive Deionization”. In: *Desalination* 249 (2009), pp. 217–223.
- [25] Norbert Brunner et al. “Water Sector Reform Policy of India: Experiences from Case Studies in Maharashtra”. In: *Journal of Policy Modeling* 32 (2010), pp. 544–561.

- [26] D. D. Caudle et al. *Electrochemical demineralization of water with carbon electrodes*. Research and Development Progress Report. Oklahoma University Research Institute, 1966.
- [27] Bhuvan: National Remote Sensing Centre. *CartoDEM Digital Elevation Maps E43B, E43H*. URL: <http://bhuvan.nrsc.gov.in/data/download/index.php?c=s&s=C1&p=cdv1&g=>.
- [28] David Leonard Chapman. "LI. A Contribution to the Theory of Electrocapillarity". In: *Philosophical Magazine*. 6th ser. 25 (148 1913), pp. 475–481.
- [29] Raghavendra Chattopadhyay and Esther Duflo. "Women as Policy Makers: Evidence from a Randomized Policy Experiment in India". In: *Econometrica* 72.5 (2004), pp. 1409–1443.
- [30] Izaak Cohen et al. "Long Term Stability of Capacitive De-ionization Processes for Water Desalination: The Challenge of Positive Electrodes Corrosion". In: *Electrochimica Acta* 106 (2013), pp. 91–100.
- [31] David Collier. "Understanding Process Tracing". In: *Political Science and Politics* 44.4 (2011), pp. 823–830.
- [32] Brian E. Conway, Erol Ayranci, and Hassan Al-Maznai. "Use of Quasi-3-Dimensional Porous Electrodes for Adsorption and Electrocatalytic Removal of Impurities from Waste-Waters". In: *Electrochimica Acta* 47 (2001), pp. 705–718.
- [33] Kai Dai et al. "NaCl Adsorption in Multi-walled Carbon Nanotubes". In: *Materials Letters* 59 (2005), pp. 1989–1992.
- [34] P. Debye. *Polar Molecules*. New York: Reinhold, 1929. Chap. VI.
- [35] *Deep Wells and Prudence. Towards Pragmatic Action for Addressing Groundwater Overexploitation in India*. The World Bank, 2010.
- [36] *Drinking Water Standards and Health Advisories*. U.S. Environmental Protection Agency, 2012.
- [37] Stephen Van Evera. *Guide to Methods for Students of Political Science*. 1997.
- [38] J. C. Farmer et al. "Capacitive Deionization of NH_4ClO_4 Solutions with Carbon Aerogel Electrodes". In: *Journal of Applied Electrochemistry* 26 (1996), pp. 1007–1018.
- [39] J. C. Farmer et al. "Capacitive Deionization with Carbon Aerogel Electrodes: Carbonate, Sulfate, and Phosphate Removal". In: *Proceedings of 17th International SAMPE Technical Conference*. 1995, pp. 294–304.
- [40] J. C. Farmer et al. "The Application of Carbon Aerogels Electrodes to Desalination and Waste Treatment". In: *Proceedings of Separation Science and Technologies Conference*. Los Angeles: AIChE, 1997.

- [41] Joseph Farmer. “Method and Apparatus for Capacitive Deionization, Electrochemical Purification, and Regeneration of Electrodes”. Patent 5425858. 1995.
- [42] Joseph C. Farmer. “Farmer and Apparatus for Capacitive Deionization and Electrochemical Purification and Regeneration of Electrodes”. Patent 5954937. 1999.
- [43] Joseph C. Farmer et al. “Capacitive Deionization of NaCl and NaNO₃ Solutions with Carbon Aerogel Electrodes”. In: *Journal of the Electrochemical Society* 143 (1 1996), pp. 159–169.
- [44] Joseph C. Farmer et al. “Electrosorption of Chromium Ions on Carbon Aerogel Electrodes as a Means of Remediating Ground Water”. In: *Energy & Fuels* 11 (1997), pp. 337–347.
- [45] J. P. Ferraris et al. “Pressure Effects on the Resistivity of Pristine and Doped Polyacetylene”. In: *Solid State Communications* 35 (1980), pp. 15–18.
- [46] R.A. Freeze and J.A. Cherry. *Groundwater*. Englewood Cliffs, NJ: Prentice-Hall, 1979.
- [47] Christopher J. Gabelich, Tri D. Tran, and I. H. ”Mel” Suffet. “Electrosorption of Inorganic Salts from Aqueous Solution Using Carbon Aerogels”. In: *Environmental Science & Technology* 36 (2002), pp. 3010–3019.
- [48] Xin Gao et al. “Surface Charge Enhanced Carbon Electrodes for Stable and Efficient Capacitive Deionization using Inverted Adsorption-Desorption Behavior”. In: *Energy & Environmental Science* 8 (2015), pp. 897–909.
- [49] Yang Gao et al. “Electrosorption Behavior of Cations with Carbon Nanotubes and Carbon Nanofibres Composite Film Electrodes”. In: *Thin Solid Films* 517 (2009), pp. 1616–1619.
- [50] Yang Gao et al. “Electrosorption of FeCl₃ Solutions with Carbon Nanotubes and Nanofibers Film Electrodes Growth on Graphite Substrates”. In: *Surface Review and Letters* 14 (6 2007), pp. 1033–1037.
- [51] Y. Gao et al. “Electrosorption of Cupric Ions from Solutions by Carbon Nanotubes and Nanofibres Film Electrodes Grown on Graphite Substrates”. In: vol. 1-3. IEEE International, 2008, pp. 242–247.
- [52] Alexander George. “Case Studies and Theory: The Method of Structured, Focused Comparison”. In: *Diplomacy: New Approaches in History, Theory, and Policy*. 1979, pp. 43–68.
- [53] Alexander L. George and Andrew Bennett. *Case Studies and Theory Development in the Social Sciences*. 2005.
- [54] Tom Gleeson et al. “Water balance of global aquifers revealed by groundwater footprint”. In: *Nature* 488 (2012), pp. 197–200.
- [55] Anatoly Golovnev and Steffen Trimper. “Analytical Solution of the Poisson-Nernst-Planck Equations in the Linear Regime at an Applied DC-Voltage”. In: *Journal of Chemical Physics* 134 (2011), p. 154902.

- [56] D. Golub, Y. Oren, and A. Soffer. “The Electrical Double Layer of Carbon and Graphite Electrodes. 4. Dependence of Carbon Electrode Dimensions and Electrical Capacity on pH”. In: *Journal of Electroanalytic Chemistry* 227 (1987), pp. 41–53.
- [57] D. Golub, A. Soffer, and Y. Oren. “The Electrical Double Layer of Carbon and Graphite Electrodes. 5. Specific Interactions with Simple Ions”. In: *Journal of Electroanalytic Chemistry* 260 (1989), pp. 383–392.
- [58] M. [Louis George] Gouy. “Sur la constitution de la charge électrique à la surface d’un électrolyte”. French. In: *Comptes Rendus de l’Académie des sciences* 149 (1909), p. 654.
- [59] M. [Louis George] Gouy. “Sur la constitution de la charge électrique à la surface d’un électrolyte”. French. In: *Journal De Physique (J. Phys. Theor. Appl.)* 1 (9 1910), p. 457.
- [60] David C. Grahame. “The Electrical Double Layer and the Theory of Electrocapillarity”. In: *Chemical Reviews* 41 (1947), pp. 441–501.
- [61] Z. J. Han et al. “Time-Dependent Electrical Double Layer with Blocking Electrode”. In: *Applied Physics Letters* 94 (2009), p. 043118.
- [62] Kelsey B. Hatzell et al. “Capacitive Deionization Concept Based on Suspension Electrodes without Ion Exchange Membranes”. In: *Electrochemistry Communications* 43 (2014), pp. 18–21.
- [63] H. von Helmholtz. “Ueber einige Gesetze der Vertheilung elektrischer Ströme in körperlichen Leitern mit Anwendung auf die thierisch-elektrischen Versuche”. German. In: *Poggendorff’s Annalen der Physik und Chemie* 89 (6 1853), pp. 211–233.
- [64] Herman Ludwig Ferdinand von Helmholtz. “Some Laws Concerning the Distribution of Electric Currents in Volume Conductors With Applications to Experiments on Animal Electricity”. Translated and abridged from the original German. In: *Proceedings of the IEEE* 92 (5 2004), pp. 868–870.
- [65] Pearson Electronics Inc. *Pearson Current Monitor Model 6027*. Datasheet. URL: <http://pearsonelectronics.com/pdf/6027.pdf>.
- [66] Sung-il Jeon et al. “Desalination via a New Membrane Capacitive Deionization Process Utilizing Flow-Electrodes”. In: *Energy & Environmental Science* 6 (2013), pp. 1471–1475.
- [67] A. M. Johnson and John Newman. “Desalting by means of porous carbon electrodes”. In: *Journal of the Electrochemical Society* 118.3 (1971), pp. 510–517.
- [68] A. M. Johnson et al. *Electrosorb process for desalting water*. Research and Development Progress Report. Office of Saline Water Research and Development, U.S. Dept. of Interior, 1970.

- [69] A. M. Johnson et al. *The Electrosorb Process for Desalting Water*. type. Van Nuys, CA: Marquardt, 1970.
- [70] Allan M. Johnson. “Electric demineralization apparatus”. Patent 3755135. 1973.
- [71] Allen M. Johnson and Richard F. Gillman. “Demineralizing Process and Apparatus”. Patent 3515664. 1970.
- [72] Hae-Hyun Jung et al. “Capacitive Deionization Characteristics of Nanostructured Carbon Aerogels Synthesized via Ambient Drying”. In: *Desalination* 216 (2007), pp. 377–385.
- [73] James L. Kaschmitter, Steven T. Mayer, and Richard W. Pekala. “Supercapacitors Based on Carbon Foams”. Patent 5260855. 1995.
- [74] Yu-Jin Kim and Jae-Hwan Choi. “Improvement of Desalination Efficiency in Capacitive Deionization Using a Carbon Electrode Coated with an Ion-Exchange Polymer”. In: *Water Research* 44 (2010), pp. 990–996.
- [75] Elizabeth Kleemeier. “The Impact of Participation on Sustainability: An Analysis of the Malawi Rural Piped Water Scheme”. In: *World Development* 28.5 (2000), pp. 929–944.
- [76] Jae-Bong Lee et al. “Desalination of a Thermal Power Plant Wastewater by Membrane Capacitive Deionization”. In: *Desalination* 196 (2006), pp. 125–134.
- [77] Philip Lessner and John Newman. “Hydrodynamics and Mass Transfer in a Porous-Wall Channel”. In: *Journal of the Electrochemical Society* 131.8 (1984), pp. 1828–1831.
- [78] C. B. Li and S. P. Friedman. “An Apparatus for Measuring the Wien Effect in Suspensions”. In: *Colloids and Surfaces A* 222 (2003), pp. 133–140.
- [79] Haibo Li et al. “Electrosorption Behavior of Graphene in NaCl Solutions”. In: *Journal of Materials Chemistry* 19 (2009), pp. 6773–6779.
- [80] Haibo Li et al. “Electrosorptive Desalination by Carbon Nanotubes and Nanofibres Electrodes and Ion-Exchange Membranes”. In: *Water Research* 42 (2008), pp. 4923–4928.
- [81] Lixia Li et al. “Ordered Mesoporous Carbons Synthesized by a Modified Sol-gel Process for Electrosorptive Removal of Sodium Chloride”. In: *Carbon* 47 (2009), pp. 775–781.
- [82] H. Lockwood and S. Smits. *Lessons for Rural Water Supply: Moving towards a Service Delivery Approach*. IRC International Water and Sanitation Centre, 2011.
- [83] S. Manikutty. “Community Participation: So What? Evidence from a Comparative Study of Two Rural Water Supply and Sanitation Projects in India”. In: *Development Policy Review* 15 (1997), pp. 115–140.

- [84] S. Manikutty. “Community Participation: Lessons from Experiences in Five Water and Sanitation Projects in India”. In: *Development Policy Review* 16 (1998), pp. 373–404.
- [85] Ghazala Mansuri and Vijayendra Rao. *Localizing Development: Does Participation Work?* Policy Research Report. The World Bank, 2013.
- [86] Sara Marks, Kyle Onda, and Jennifer Davis. “Does Sense of Ownership Matter for Rural Water System Sustainability? Evidence from Kenya”. In: *Journal of Water, Sanitation and Hygiene for Development* 3.2 (2013), pp. 122–133.
- [87] Michael Matlosz and John Newman. “Experimental Investigation of a Porous Carbon Electrode for the Removal of Mercury from Contaminated Brine”. In: *Journal of the Electrochemical Society* 133.9 (1986), pp. 1850–1858.
- [88] S. T. Mayer, R. W. Pekala, and J. L. Kaschnmitter. “The Aerocapacitor: An Electrochemical Double-Layer Energy-Storage Device”. In: *Journal of the Electrochemical Society* 140.2 (1993), pp. 446–451.
- [89] Steven T. Mayer, James L. Kaschnmitter, and Richard W. Pekala. “Carbon Aerogels for Direct Energy Conversion”. Patent 5601938. 1997.
- [90] John Stewart Mill. *A System of Logic*. 1847.
- [91] R. Morrow, D. R. McKenzie, and M. M. M. Bilek. “The Time-Dependent Development of Electric Double-Layers in Saline Solutions”. In: *Journal of Physics. D: Applied Physics* 39 (2006), pp. 937–943.
- [92] G. W. Murphy and D. D. Caudle. “Mathematical theory of electrochemical demineralization in flowing systems”. In: *Electrochimica Acta* 12 (1967), pp. 1655–1664.
- [93] G. W. Murphy, J. L. Cooper, and J. A. Hunter. *Activated carbon used as electrodes in electrochemical demineralization of saline water*. U.S. Dept of the Interior, 1969.
- [94] Deepa Narayan. *The Contribution of People’s Participation: Evidence from 121 Rural Water Supply Projects*. The World Bank, 1995.
- [95] John Newman and Karen E. Thomas-Alyea. *Electrochemical Systems*. 3rd. Wiley, 2004.
- [96] Lars Onsager and Shoon Kyung Kim. “Wien Effect in Simple Strong Electrolytes”. In: *Journal of Physical Chemistry* 61 (2 1957), pp. 198–215.
- [97] Y. Oren and A. Soffer. “Electrochemical Parametric Pumping”. In: *Journal of the Electrochemical Society* 125.6 (1978), pp. 869–875.
- [98] Y. Oren and A. Soffer. “Equilibrium Properties of an Electrochemical Desalination Unit Composed of 2 High Surface-Area Carbon Electrodes”. In: *Journal of the Electrochemical Society* 126 (8 1979), p. C330.
- [99] Y. Oren and A. Soffer. “The Electrical Double Layer of Carbon and Graphite Electrodes. 2. Fast and Slow Charging Processes”. In: *Journal of Electroanalytic Chemistry* 186 (1985), pp. 63–77.

- [100] Y. Oren and A. Soffer. “Water Desalination as a Test Model for Electroadsorption Parametric Pumping”. In: *Journal of the Electrochemical Society* 124.3 (1977), p. C121.
- [101] Y. Oren and A. Soffer. “Water Desalting by Means of Electrochemical Parametric Pumping. 1. The Equilibrium Properties of a Batch Unit-Cell”. In: *Journal of Applied Electrochemistry* 13 (4 1983), pp. 473–487.
- [102] Y. Oren and A. Soffer. “Water Desalting by Means of Electrochemical Parametric Pumping. 2. Separation Properties of a Multistage Column”. In: *Journal of Applied Electrochemistry* 13 (4 1983), pp. 489–505.
- [103] Y. Oren, H. Tobias, and A. Soffer. “Removal of Bacteria from Water by Electroadsorption on Porous Carbon Electrodes”. In: *Bioelectrochemistry and Bioenergetics* 11 (1983), pp. 347–351.
- [104] Y. Oren, H. Tobias, and A. Soffer. “The Electrical Double Layer of Carbon and Graphite Electrodes. 1. Dependence on Electrolyte Type and Concentration”. In: *Journal of Electroanalytic Chemistry* 162 (1984), pp. 87–99.
- [105] Yoram Oren. “Capacitive deionization (CDI) for desalination and water treatment - past, present and future (a review)”. In: *Desalination* 228 (2008), pp. 10–29.
- [106] Yoram Oren and Abraham Soffer. “The Electrical Double Layer of Carbon and Graphite Electrodes. 3. Charge and Dimensional Changes at Wide Potential Range”. In: *Journal of Electroanalytic Chemistry* 206 (1986), pp. 101–114.
- [107] Kwang-Kyu Park et al. “Development of a Carbon Sheet Electrode for Electrosorption Desalination”. In: *Desalination* 206 (1-3 2007), pp. 86–91.
- [108] R. W. Pekala. “Organic Aerogels from the Polycondensation of Resorcinol with Formaldehyde”. In: *Journal of Material Science* 24 (9 1989), pp. 3221–3227.
- [109] R. W. Pekala and C. T. Alviso. “Carbon Aerogels and Xerogels”. In: *Materials Research Society Symposium Proceedings*. Ed. by C. L. Renschler, J. J. Pouch, and D. M. Cox. Vol. 270. 1992, pp. 3–14.
- [110] R. W. Pekala et al. “Carbon Aerogels for Electrochemical Applications”. In: *Journal of Non-Crystalline Solids* 225 (1998), pp. 74–80.
- [111] R. W. Pekala et al. “Structure and Performance of Carbon Aerogel Electrodes”. In: *Novel Forms of Carbon II*. Materials Research Society Symposium. Ed. by C. L. Renschler et al. Vol. 349. 79. 1994.
- [112] Richard W. Pekala. “Melamine-formaldehyde Aerogels”. Patent 5081163. 1992.
- [113] S. Porada et al. “Direct Prediction of the Desalination Performance of Porous Carbon Electrodes for Capacitive Deionization”. In: *Energy & Environmental Science* 6 (2013), pp. 3700–3712.
- [114] S. Porada et al. “Review on the science and technology of water desalination by capacitive deionization”. In: *Progress in Materials Science* 58 (2013), pp. 1388–1442.

- [115] S. Porada et al. “Water Desalination with Wires”. In: *Journal of Physical Chemistry Letters* 3 (2012), pp. 1613–1618.
- [116] National Rural Drinking Water Programme. *Maharashtra Database B: B11-Rural Population*. URL: http://indiawater.gov.in/IMISReports/Reports/BasicInformation/rpt_RWS_RuralPopulation_D.aspx?Rep=1&RP=Y.
- [117] National Rural Drinking Water Programme. *Maharashtra Database B: B15-Complete List of Schemes*. URL: http://indiawater.gov.in/imisreports/Reports/BasicInformation/rpt_RWS_SchemeStatusSelection.aspx?Rep=0&RP=Y.
- [118] WWAP (United Nations World Water Assessment Programme). *The United Nations Development Report 2015. Water for a Sustainable World*. UNESCO, 2015.
- [119] Linda Stalker Prokopy. “The Relationship Between Participation and Project Outcomes: Evidence from Rural Water Supply Projects in India”. In: *World Development* 33.11 (2005), pp. 1801–1819.
- [120] Government of Raigad District. *Raste Vikaas Yojna Karjat Taluka (Road Development Map for Karjat Subdistrict)*. Marathi. URL: <http://raigad.gov.in/Images/karjat%20taluka.jpg>.
- [121] Ben Ramalingham. *Aid on the Edge of Chaos*. Oxford, UK: Oxford Univeristy Press, 2013.
- [122] Adam Rausch, Ashok Gadgil, and Robert Kosteki. “Devices and Methods for Removing Dissolved Ions from Water Using a Voltage-driven Charge Pulse”. Patent Application PCT/US15/54091. 2015.
- [123] *R&D100 Conference: Award Winners and Finalists*. 1991. URL: www.rd100conference.com/awards/winners-finalists/year/1991/.
- [124] Tim Risch and John Newman. “A Theoretical Comparison of Flow-Through and Flow-By Porous Electrodes at the Limiting Current”. In: *Journal of the Electrochemical Society* 131.11 (1984), pp. 2551–2556.
- [125] Min-Woong Ryoo, Jong-Ho Kim, and Gon Seo. “Role of Titania Incorporated on Activated Carbon Cloth for Capacitive Deionization of NaCl solution”. In: *Journal of Colloid and Interface Science* 264 (2003), pp. 414–419.
- [126] Min-Woong Ryoo and Gon Seo. “Improvement in Capacitive Deinization Function of Activated Carbon Cloth by Titania Modification”. In: *Water Research* 37 (2003), pp. 1527–1534.
- [127] Jennifer Sara and Travis Katz. *Making Rural Water Supply Sustainable: Report on the Impact of Project Rules*. UN Development Programme - World Bank Water and Sanitation Program, 1997.
- [128] Ketki Sharma et al. “Enhancement of Electrosorption Rates Using Low-Amplitude, High-Frequency, Pulsed Electric Potential”. In: *Sorption and Purification Technology* 129 (2014), pp. 18–24.

- [129] Otto Stern. “Zur Theorie der Elektrolytischen Doppelschicht”. German. In: *Zeitschrift für Elektrochemie* 30 (1924), pp. 508–516.
- [130] M. E. Suss et al. “Water Desalination /textitvia Capacitive Deionization: What is it and what can we expect from it?” In: *Energy & Environmental Science* 8 (2015), pp. 2296–2319.
- [131] Matthew E. Suss et al. “Capacitive Desalination with Flow-Through Electrodes”. In: *Energy & Environmental Science* 5 (2012), pp. 9511–9519.
- [132] Matthew E. Suss et al. “Impedance-Based Study of Capacitive Porous Carbon Electrodes with Hierarchical and Bimodal Porosity”. In: *Journal of Power Sources* 241 (2013), pp. 266–273.
- [133] T. D. Tran et al. “Effects of Processing Condidtions on the Physical and Electrochemical Properties of Carbon Aerogel Composites”. In: *Materials for Electrochemcial Storage and Conversion II—Batteries, Capacitors, and Fuel Cells*. Vol. 496. Warrendale, PA: Materials Research Society, 1998.
- [134] Tri D. Tran, Joseph C. Farmer, and Laura Murguia. “Method and Apparatus for Capacitive Deionization and Electrochemical Purification and Regeneration of Electrodes”. Patent 6309532. 2001.
- [135] Tri D. Tran and David J. Lenz. “Alternating-Polarity Operation for Complete Regeneration of Electrochemical Deionization System”. Patent 6346197. 2002.
- [136] Tri D. Tran and David J. Lenz. “Alternating-Polarity Operation for Complete Regeneration of Electrochemical Deionization System”. Patent 6761809. 2004.
- [137] Tri D. Tran and David J. Lenz. “Alternating-Polarity Operation for Complete Regeneration of Electrochemical Deionization System”. Patent 7138042. 2006.
- [138] Hui Wang et al. “Three-Dimensional Macroporous Graphene Architectures as High Performance Electrodes for Capacitive Deionization”. In: *Journal of Materials Chemistry A* 1 (2013), pp. 11778–11789.
- [139] X. Z. Wang et al. “Electrosorption of Ions from Aqueous Solutions with Carbon Nanotubes and Nanofibers Composite Film Electrodes”. In: *Applied Physics Letters* 89 (5 2006).
- [140] Frank van Weert, Jac van der Gun, and Josef Reckman. *Global Overview of Saline Groundwater Occurence and Genesis*. International Groundwater Resources Assessment Centre, 2009.
- [141] T. J. Welgemoed and C. F. Schutte. “Capacitive Deionization TechnologyTM: An Alternative Desalination Solution”. In: *Desalination* 183 (2005), pp. 327–340.
- [142] Von Max Wien. “2. Über den Spannungseffekt der elektrolytischen Leitfähigkeit in sehr starken Feldern (The tension effect of electrolytic conductivity in very strong fields)”. German. In: *Annalen der Physik* 1 (3 1929), pp. 400–416.

- [143] Von Max Wien. “Über eine Abweichung vom Ohmschen Gesetze bei Elektrolyten (On a discrepancy from Ohm’s law on electrolytes)”. German. In: *Annalen der Physik* 83 (11 1927), pp. 327–361.
- [144] Pei Xu et al. “Treatment of Brackish produced Water Using Carbon Aerogel-Based Capacitive Deionization Technology”. In: *Water Research* 42 (2008), pp. 2605–2617.
- [145] Juan Yang, Linda Zou, and Nmita Roy Choudhury. “Ion-Selective Carbon Nanotube Electrodes in Capacitive Deionisation”. In: *Electrochimica Acta* 91 (2013), pp. 11–19.
- [146] Kun-Lin Yang et al. “Electrosorption of Ions from Aqueous Solutions by Carbon Aerogel: An Electrical Double-Layer Model”. In: *Langmuir* 17 (2001), pp. 1961–1969.
- [147] Tung-Yu Ying et al. “Electrosorption of Ions from Aqueous Solutions by Nanostructured Carbon Aerogel”. In: *Journal of Colloid and Interface Sciences* 250 (2002), pp. 18–27.
- [148] Yankun Zhan et al. “Carbon Nanotube-Chitosan Composite Electrodes for Electrochemical Removal of Cu(II) Ions”. In: *Journal of Alloys and Compounds* 509 (2011), pp. 5667–5671.
- [149] R. Zhao, P. M. Biesheuvel, and A. van der Wal. “Energy Consumption and Constant Current Operation in Membrane Capacitive Deionization”. In: *Energy & Environmental Science* 5.11 (2012), p. 9520.
- [150] Linda Zou, Gayle Morris, and Daoduo Qi. “Using Activated Carbon Electrode in Electrosorptive Dionisation of Brackish Water”. In: *Desalination* 225 (2008), pp. 329–340.
- [151] Linda Zou et al. “Using Mesoporous Carbon Electrodes for Brackish Water Desalination”. In: *Water Research* 42 (2008), pp. 2340–2348.