ATTACHMENT 12: CDISCO Description and Sensitivity to Input Parameters

INTRODUCTION

The A11. ISCO Spreadsheet Design Tool (a.k.a. Conceptual Design for ISCO or CDISCO) was developed with support from ESTCP under Projects ER-0625 and ER-0623. It is intended to assist with the conceptual design of injection systems for (ISCO) using permanganate. The tool "models" the transport and degradation of permanganate during the injection period and for a short time after injection. It does not model downgradient advective drift of oxidant after the injection period. Additionally, it does not model the contaminant concentrations because of the complexities of the various phases of the contaminants that may be present.

CDISCO DESCRIPTION

CDISCO consists of two major components.

1 A numerical model for simulating the 1D radial transport and consumption of permanganate following injection into a single well. This model allows users to easily evaluate the effect of aquifer parameters (effective thickness, permeability, total natural oxidant demand (NOD), NOD kinetics) and injection conditions (permanganate concentration, injection flow rate and duration) on the effective radius of influence (ROI) of an individual injection well, and required well spacing based on a user defined overlap factor.

2 A procedure which allows designers to quickly calculate a preliminary and partial cost estimate to evaluate the effect of injection conditions (permanganate concentration, injection flow rate and duration) on injection costs. Input parameters include reinjection frequency, injection fixed costs and unit costs for injection point installation, chemical reagents, and labor for injection.

The tool allows the user to enter site specific data or assumed values of key site and design parameters. Design parameters, such as injection method (well versus direct push point), duration of oxidant injection, mass of oxidant injected, and volume of oxidant solution injected can be varied and the tool will estimate well/point spacing and generate a preliminary cost estimate. This allows a preliminary design optimization to be performed by quickly executing various model runs with varied inputs and comparing the resulting cost. Graphical representations of the effect of the design parameters on project costs are generated.

The tool was designed primarily for permanganate; however, with modification to the kinetic input parameters, it is possible that CDISCO can approximate pseudo first-order kinetics applied to other oxidant persistence. Possible means to accomplish this, along with limitations of doing so, are discussed further below.

Permanganate Transport Model

Advective-dispersive transport of permanganate away from a central injection well is represented in CDISCO as flow through a series of continuously stirred tank reactors (CSTR). Numerical dispersion associated with evaluation of the time derivative is negligible (Van Genuchten and Wierenga, 1974) so longitudinal dispersion is simulated by setting the length of each reactor to two times the longitudinal dispersivity, which is selected by the user. Permanganate consumption by the target contaminant is simulated as an instantaneous reaction. The mass of permanganate consumed by the target contaminant is typically small compared to the mass injected, unless very high (NAPL levels) of contaminant are assumed to be present.

The NOD is assumed to be composed of two fractions – NOD_I which reacts instantaneously with permanganate and NOD_S which reacts more slowly. After the instantaneous NOD is satisfied, the rate of permanganate (*M*) loss with time (*dM / dt*) is proportional to the oxidant concentration times the NOD_S concentration.

$$\frac{dM}{dt} = -k_S NOD_S M \rho_B / n$$

where k_s is the 2nd order slow NOD consumption rate, ρ_B is density, and *n* is porosity. Separation of the NOD into instantaneous (or very fast) and slow fractions is consistent with experimental results from Mumford et al. (2005), Urynowicz et al. (2008), Xu (2006), and others.

The CSTR model is implemented within a Microsoft Excel spreadsheet. The user first enters information on aquifer characteristics (porosity, hydraulic conductivity, injection interval, NOD, contaminant concentrations, etc.), injection conditions (permanganate injection concentration, flow rate and duration), and target conditions (minimum oxidant concentration and duration to calculate ROI). Based on this information, the spreadsheet establishes a series of reactors. The volume of each reactor increases outward as injected water migrates radially outward from the injection well, simulating the reduced velocity as flow radiates out from the injection well. The actual computations to determine spatial and temporal variations in permanganate concentration are performed in a Visual Basic Macro embedded in the CDISCO spreadsheet. The time step for each computation is automatically determined within Excel to minimize computational error. The macro calculates the permanganate concentration in each reactor sequentially and then at each time step. It builds a table of permanganate concentrations with distance and time. It then selects the radius of influence (ROI) from a user defined minimum concentration of permanganate at a minimum time period after the injection. For example, the user may decide that they want a minimum of 100 mg/l of permanganate to persist for 7 days after injection. CDISCO will select the radius of influence that meets this criteria from the table of oxidant concentrations versus time and distances.

Figure A12-1 shows typical output from the permanganate transport model simulations. The graph shows the variation in permanganate concentration versus distance at several different times (15, 30, 45 and 60 days for this simulation). The table at the bottom shows input parameters and computed ROI for a series of prior simulations.

Cost Estimating Procedure

The cost estimating procedure generates a preliminary cost estimate for the injection based on the user specified treatment area dimensions, injection well ROI overlap (%), number of injection events planned, fixed cost, and unit costs for injection point installation, chemical reagents, and labor for injection. Two injection approaches are possible – injection through direct push rods or through wells. Cost factors are included for mobilization, labor, materials, equipment rental, travel, and subcontractor costs.

Figure A12-2 shows typical output from the cost estimating procedure comparing preliminary estimates for several different design alternatives. Aquifer parameters, treatment area dimensions (100 ft x 100 ft), ROI overlap (25%), time to calculate ROI (30 days), minimum oxidant concentration (50mg/L), and number of injection events (5) are constant for all the alternatives in this example. Alternative 1 has a relatively high cost because the short injection duration (3 days) required a large number of injection points. Alternative 5 has a lower cost because the longer injection duration (10 days) and higher oxidant concentration (20,000 mg/L $KMnO_4$) reduced number of injection points required. Multiple similar runs can be performed to allow the user to optimize the design for the given site conditions, and also to evaluate the sensitivity of the cost to site conditions (e.g. the NOD). If the costs are very sensitive to NOD, and there is significant uncertainty in the actual NOD, more field data may need to be collected.

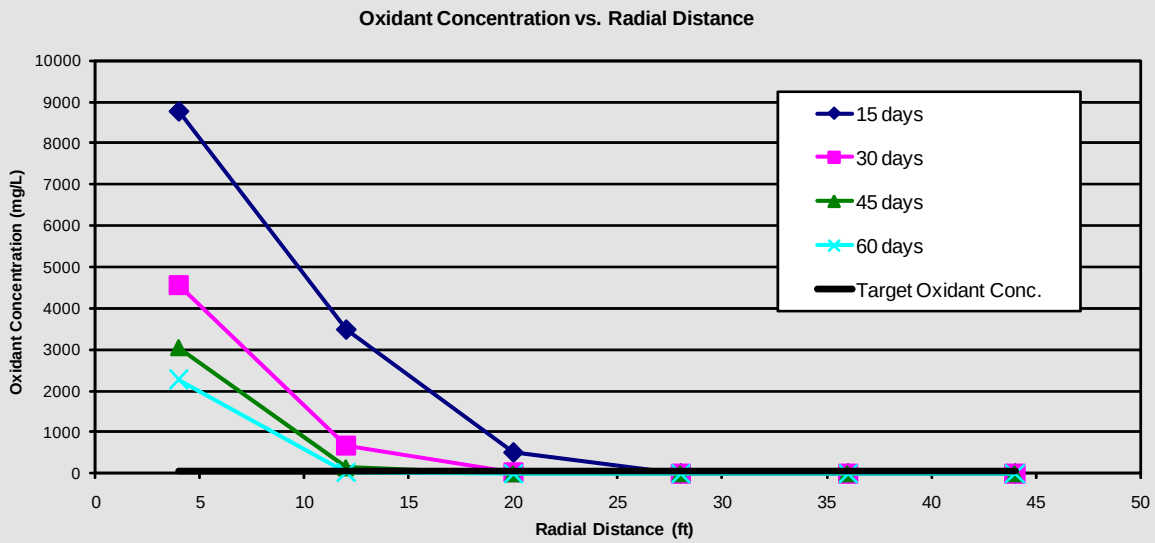

Selected	Run Number	Injection Duration	Aquifer Thickness	Thickness of Mobile Zone	NOD Fraction Instantaneous	NOD (g/kg)	Slow NOD Rate	Injection Oxidant Conc (mg/L)	Injection Rate gal/Day	ROI	Minimum Oxidant Conc to Calc ROI	Target Number of Days to Calc ROI
✔	1	3	20	10	0.1	5	0.001	10,000	3,000	10.41	50	30
✔	2	5	20	10	0.1	5	0.001	10,000	3,000	14.46	50	30
✔	3	10	20	10	0.1	5	0.001	10,000	3,000	21.25	50	30
✔	4	10	20	10	0.1	5	0.001	5,000	3,000	15.56	50	30
✔	5	10	20	10	0.1	5	0.001	20,000	3,000	27.23	50	30
✔	6	5	20	10	0.1	5	0.001	20,000	3,000	19.64	50	30

Figure A12-1: Typical Output from Permanganate Transport Model.

Run	1	2	3	4	5	6
Total Fixed Costs (injection)	$94,800	$94,800	$94,800	$94,800	$94,800	$94,800
Total Well Installation Costs	$85,667	$47,700	$25,367	$41,000	$18,667	$29,833
Total Injection Costs	$478,800	$410,400	$364,800	$684,000	$228,000	$228,000
Total Oxidant Cost	$378,547	$324,469	$288,417	$270,391	$360,521	$360,521
Total Installation and Injection Costs	$1,037,814	$877,369	$773,384	$1,090,191	$701,988	$713,155
Number of probes or wells required	35	18	8	15	5	10
NOD (g/kg)	5	5	5	5	5	5
Injection Oxidant Concentration	10000	10000	10000	5000	20000	20000
Injection Oxidant Mass (lbs)	26288	22533	20029	18777	25036	25036
Injection Duration (days)	3	5	10	10	10	5
Volume Injected per Day (gal/d)	3000	3000	3000	3000	3000	3000
Thickness of Mobile/Target Thickness	0.5	0.5	0.5	0.5	0.5	0.5

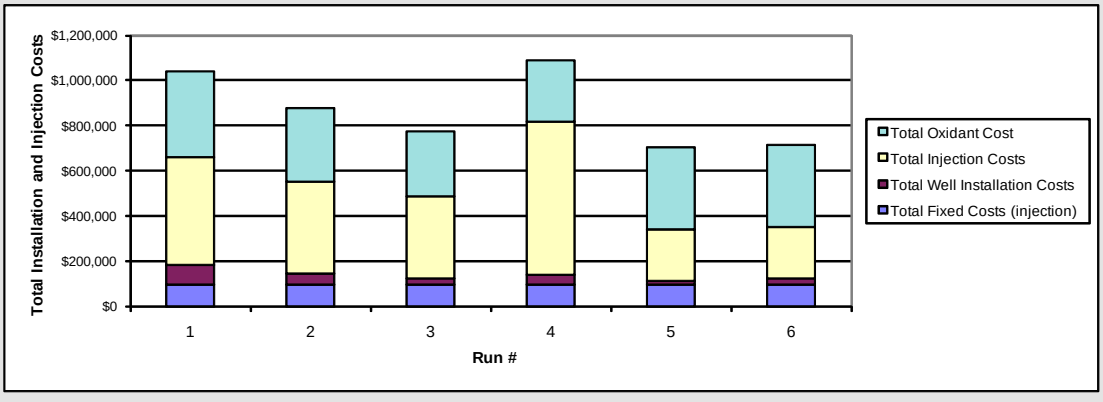

Figure A12-2. Typical Output from Injection Scenario Cost Comparison.

Application of CDISCO to Persulfate and CHP

CDISCO was initially developed to assist in the design of permanganate injection systems. However, it may be possible to simulate some elements of CHP and persulfate distribution with appropriate modifications to the CDISCO input parameters. The decomposition of hydrogen peroxide in well-mixed batch systems is often observed to follow pseudo first-order kinetics. CDISCO can be forced to simulate pseudo first-order decay of the oxidant by setting the contaminant concentration and instantaneous NOD equal to zero and setting the slow NOD (NOD_S) much greater than the maximum oxidant concentration. Under these conditions, the effective 1^{st} order decay rate will be equal to the 2^{nd} order slow NOD decay rate times NOD_S. Where persulfate loss is 1^{st} order, this same approach could potentially be applied.

One major limitation on the use of CDISCO for simulating persulfate and CHP distribution is that it does not consider the transport of an activator (e.g., iron, base, another oxidant, heat, etc.). Since the subsurface transport of the activator may be very different than the primary oxidant, the pseudo 1^{st} order rate may vary spatially and with time. Under these conditions, it would not be appropriate to use CDISCO to model persulfate and CHP distribution. In general, the only situation where CDISCO may approximate the kinetics of oxidant persistence is where the activator and oxidant are mixed simultaneously at the injection point, and then the assumption must be made that the first-order rate stays constant from that time onward.

Potential users should also be aware that the reactive transport of persulfate and CHP is not well documented and there are no good datasets suitable for verifying CDISCO with these oxidants. Thus, in most cases, it will not be practical to use CDISCO for predicting persulfate or CHP transport and persistence. Any user attempting to use the CDISCO tool for oxidants other than permanganate should be forewarned that there is considerable uncertainty in how accurately the tool output will match real field processes, and caution is advised. Users should have a good understanding of ISCO before using this tool.

SENSITIVITY ANALYSIS

This parameter sensitivity analysis was performed to evaluate the output of the A11. ISCO Spreadsheet Design Tool (also called Conceptual Design for ISCO or CDISCO) under a number of input conditions. The output was only evaluated in terms of the ROI achieved by the oxidant and delivery system that is estimated by the tool and not in terms of cost. It is not meant to be an exhaustive sensitivity analysis nor is it designed to validate the tool.

It should be remembered that this tool is not designed to be an all-encompassing predictor of ISCO performance. Rather, it is designed to allow the user to explore the relationship of expected performance (in terms of ROI and possible cost) with varying input and design parameters. Its primary purpose is to allow the user to make more informed decisions regarding data collection for ISCO systems, by highlighting those parameters that are most likely to impact design and cost certainty. It also allows the user to explore the impacts of design parameters (e.g., injection flows and concentrations) on the injection costs, thereby allowing a more optimized design.

Highly Sensitive Parameters

The most sensitive model input parameters under the base conditions tested are: total NOD concentration (between 0 and 3 g/kg); aquifer mobile zone thickness; injection rate (especially for values below 1,000 gal/day); and injected oxidant concentration. It is difficult to directly compare raw output values due to differences in units among parameters, but qualitative comparisons are made below (see the final discussions for quantitative output plots).

Varying NOD concentrations between 0 and 3 g/kg produced ROI output variation of up to 30 feet, indicating that this model is most sensitive to this input. This sensitivity is as expected based on experience in the field, and the model equations and formulations were designed to primarily account for

consumption of oxidant by NOD. Likewise, the similarly sensitive effect of varying injected oxidant concentrations and injection rates should be intuitive to users and it is likely that users will intentionally explore the bounds of this parameter on their own.

Model output is strongly affected by the mobile zone thicknesses. This parameter defines the volume of the aquifer that must be contacted before the oxidant solution can move any further away from the injection point. This parameter will always be a fraction of the total aquifer thickness, and it is difficult to estimate. Thus, the effect of mobile zone thickness should be evaluated by the user.

The model is also very sensitive to high values of initial contaminant concentration (i.e., NAPL-phase values), but this should only be a concern to users trying to simulate pure-phase scenarios. The provided guidance for the initial contaminant concentration informs users that raising the input value may significantly affect output, so this sensitivity is less of a concern to the general user.

Base Case Conditions

CDISCO was run under base case conditions to provide a reference value for the model output (calculated ROI), to which all subsequent runs are then compared to assess model sensitivity. The base case conditions for each permanganate/trichloroethene model run are shown below in Table A12-1. Note that two base NOD conditions are used for all other parameter sensitivity simulations: a relatively low value of 0.5 g/kg NOD, and a moderately high value of 3 g/kg.

Table A12-1. Amended Base Case Conditions

Variable Parameters		
Input Parameter	**Base Condition**	**Variation**
ROI target # of days	7 days	2 – 50 days
Aquifer thickness	20 feet	minimum – 50 feet
Mobile zone thickness	10 feet	minimum – maximum
Porosity	0.2	minimum – 0.6
Longitudinal dispersivity	0.5 feet	0.5 – 10 feet
Bulk density	1.6 kg/L	1.2 – 2.0 kg/L
NOD	0.5 g/kg; 3 g/kg	0 – 20 g/kg
NOD fraction instantaneous	0.2	0 – 1
Slow NOD reaction rate	0.01 L/mmol/day	1×10^{-5} – 1 L/mmol/day
Initial contaminant concentration	0.06 mg/L	6×10^{-5} – 6×10^{4} mg/L
Contaminant retardation	1.5	1 – 1×10^{-5}
Injected oxidant concentration	10,000 mg/L	1,000 – 40,000 mg/L
Injection rate	15,000 gal/day	200 – 15,000 gal/day
Fixed-Value Parameters		
Model duration	60 days	-
Time step	1 day	-
Model length	40 feet	-
Minimum ROI oxidant value	200 mg/L	-
Aquifer depths	30 – 50 feet bgs	-
Injection duration	3 days	-

For each sensitivity simulation, one input parameter was varied over a realistic range while all other parameters were held constant. It is recognized that this type of single-variant sensitivity analysis may not capture the co-relations among all parameters, but it does provide a gross indication of which parameters require the most confidence.

Curves of oxidant concentrations versus distance are shown in Figures A12-3 and A12-4 for four time steps under the base case conditions. The reference target number of days for ROI calculation, 7 days, is shown in red in these figures. It should be noted that the oxidant concentration does not drop further after about 14 days in this simulations because the NOD has been consumed within the area contacted.

Since the model does not simulate the advective migration with natural groundwater flow (drift) of oxidant, the results are not completely accurate for long time periods and/or for fast moving groundwater systems.

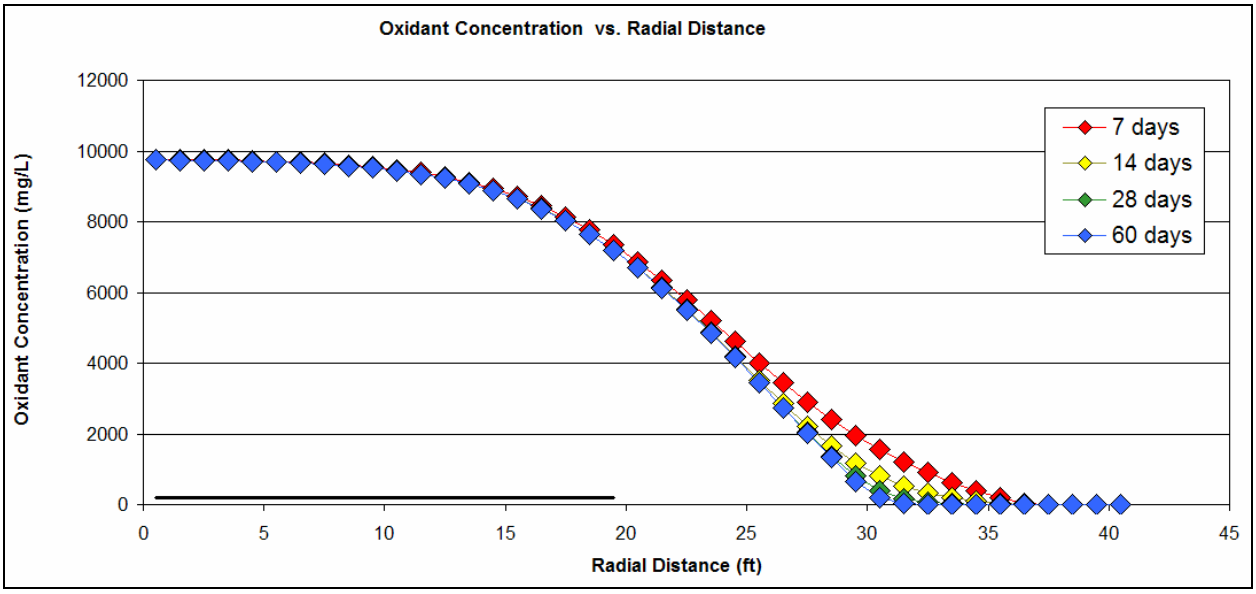

Figure A12-3: Base Condition Results, 0.5 g/kg NOD; 7-day ROI = 35.52 feet.

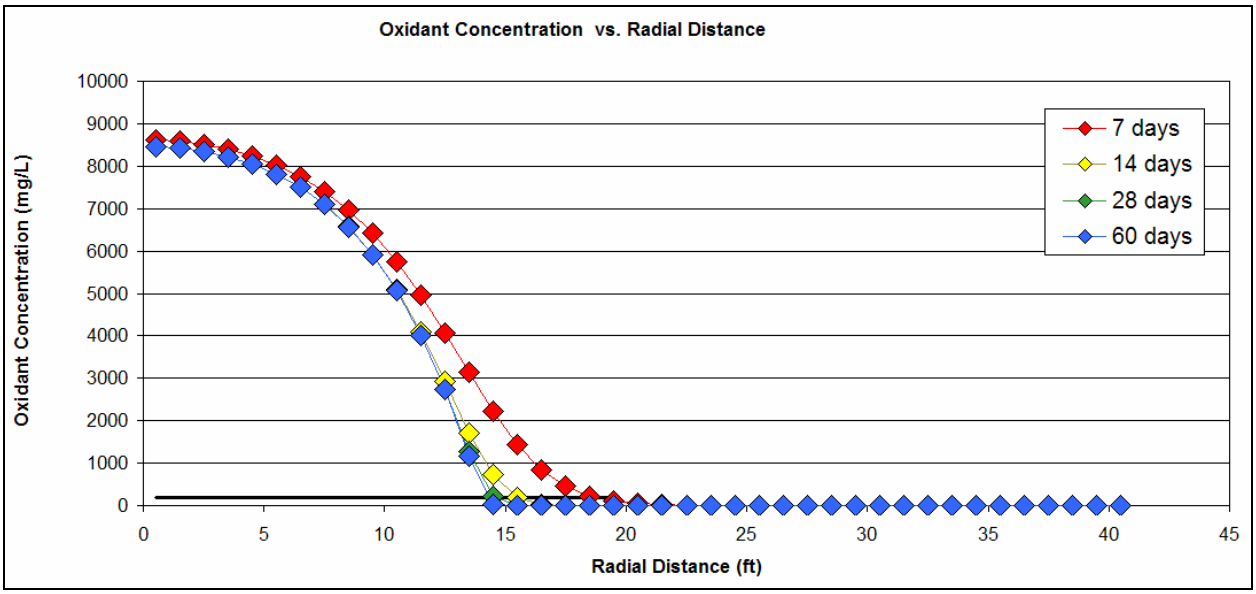

Figure A12-4: Base Condition Results, 3 g/kg NOD; 7-day ROI = 18.70 feet.

INDIVIDUAL PARAMETER SENSITIVITY

Target Number of Days for ROI

This parameter defines the minimum number of days that the user would like to have oxidant present in the subsurface, at the target concentration (selected as 200 mg/L in this case). For this simulation, the value was varied from 2 days (the minimum allowable value) to 50 days. The calculated ROI increases until the 4th day, and then begins to decline with time (Figure A12-5). It should be noted that the duration of oxidant injection was 3 days, so that the rise in ROI up to 4 days should be expected, since oxidant is

still being delivered. The ROI then declines after oxidant injection stops, which is expected, and the decline is more pronounced for the 3 g/kg NOD case than for the lower NOD case as a result of increased slow NOD reaction and oxidant consumption.

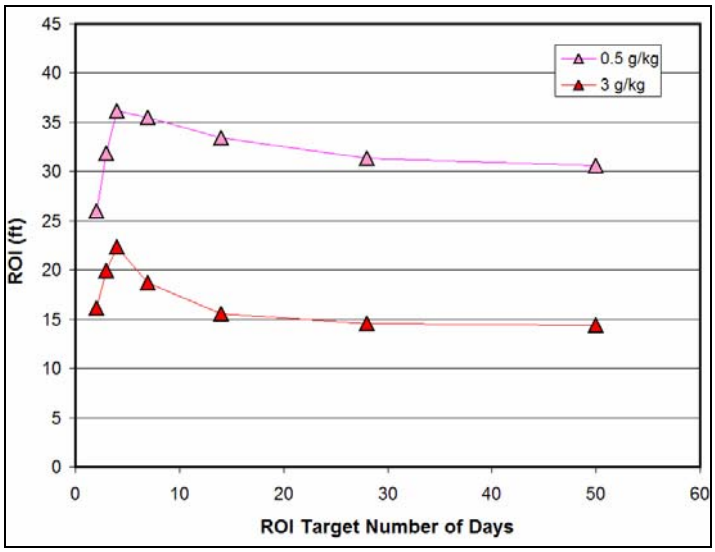

Figure A12-5: ROI Output vs. ROI Target Number of Days.

Thickness of Mobile Zone

The thickness of the mobile zone defines the fraction of the aquifer thickness that must be contacted before the oxidant solution can move any further away from the injection point. This parameter was varied over its entire possible range, from the minimum thickness value that produces an ROI within the 40-feet model domain (7.7 and 2.1 feet), to the maximum of 20 feet (100% of the available aquifer thickness). As seen for the aquifer thickness parameter, ROI gradually decreases as mobile zone thickness increases, and a straight-line relationship can be obtained by normalizing the mobile zone thickness to the cylindrical geometry (Figure A12-6). For a higher NOD value, ROI output is lower for a given mobile zone thickness.

Simplistically, having a low thickness of mobile zone is good in terms of getting larger ROI, but it may be bad in terms of terms of aquifer remediation since less of the formation will be contacted by the oxidant. If the contaminants are also only in this same mobile zone, then cleanup should be effective. But, if the contaminants have migrated (via diffusion, etc.) into the less mobile zone, then cleanup will be less effective.

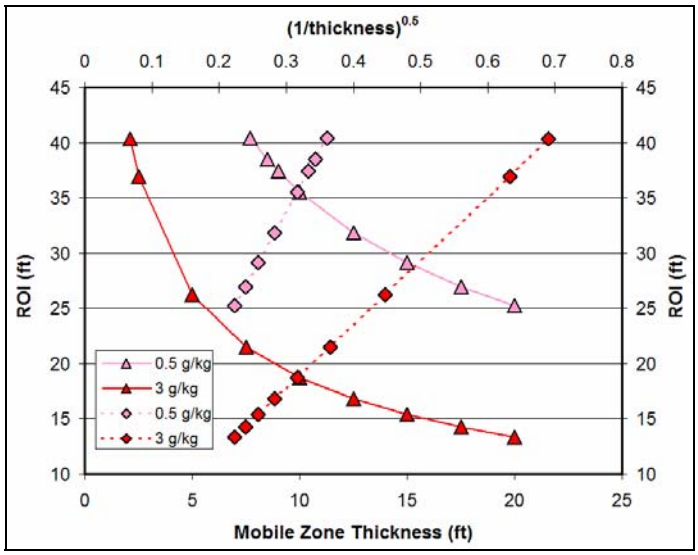

Figure A12-6: ROI Output vs. Mobile Zone Thickness (triangle symbols and solid lines) and ROI Output vs. Normalized Mobile Zone Thickness (diamond symbols and dashed lines).

Porosity

This parameter was varied from a minimum porosity that produces an ROI within the 40-feet model domain (13 and 6%) to a maximum porosity of 60%. The response of ROI to varying porosity is not constant for differing NOD cases. To further illustrate this response, a very high NOD case (10 g/kg) is also shown (dark red symbols in Figure A12-7). For low NOD cases, ROI gradually decreases with increasing porosity, much as for aquifer mobile zone thickness. When the porosity is low, a small mass of oxidant is initially transported a further distance away from the injection point. If the NOD is low, this results in a larger ROI. If the NOD is high, the small mass of oxidant is consumed with in the 7 days used to calculate the ROI. Consequently, the ROI is deemed to be small.

When the porosity is high, a large mass of oxidant is transported a shorter distance from the injection point. If the NOD is low this results in a smaller ROI. If the NOD is high, the larger mass of oxidant is not consumed within 7 days and the estimate ROI will be higher than with lower porosity. This same type of relationship might be expected in the real world.

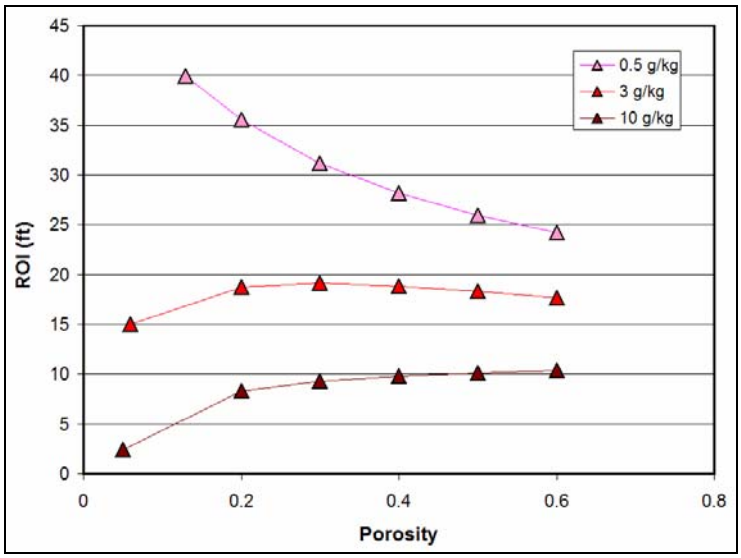

Figure A12-7: ROI Output vs. Porosity.

Longitudinal Dispersivity (or, one-half model reactor length)

The longitudinal dispersivity is used in the model to calculate the length of each reactor used in the Completely Mixed Stirred Reactor approach. The reactor length equals 2 times the longitudinal dispersivity. Consequently, it dictates the number of reactors (model length divided by reactor length). The guidance provided with CDISCO suggests at least 10 reactors be used.

The longitudinal dispersivity was varied from its base condition of 0.5 feet (resulting in reactor lengths of 1 foot and 41 reactors) to a maximum value of 10 feet (resulting in reactor lengths of 20 feet and 3 reactors). For the low NOD case, ROI values are not calculated for dispersivity values greater than 5 feet, as the estimated ROI exceeded the allowed model domain of 40 feet.

The overall trend is one of increasing ROI output with increased dispersivity, resulting in a difference of ~10 feet for ROI output (Figure A12-8). If the range of allowable dispersivity values is limited to those which permit at least ten reactors in the model domain (0.5–2 feet), the difference in ROI output is reduced to ~2 feet.

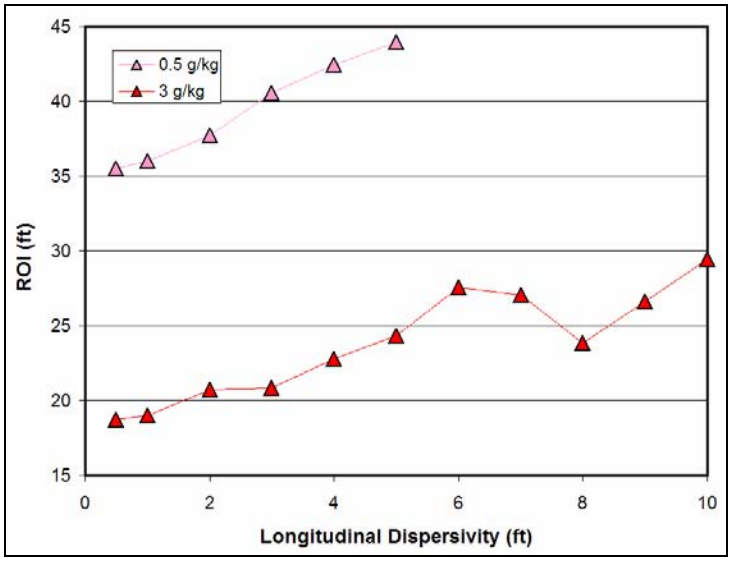

Figure A12-8: ROI Output vs. Longitudinal Dispersivity (half-reactor length).

Bulk Density

The bulk density parameter was varied from 1.2 to 2.0 kg/L. ROI values decrease slightly with increases in bulk density, but the overall change in ROI output is only 3.5 feet for the low NOD case and 5.5 feet for the moderate case. The model is fairly insensitive to this parameter within this reasonable range of bulk density values (Figure A12-9).

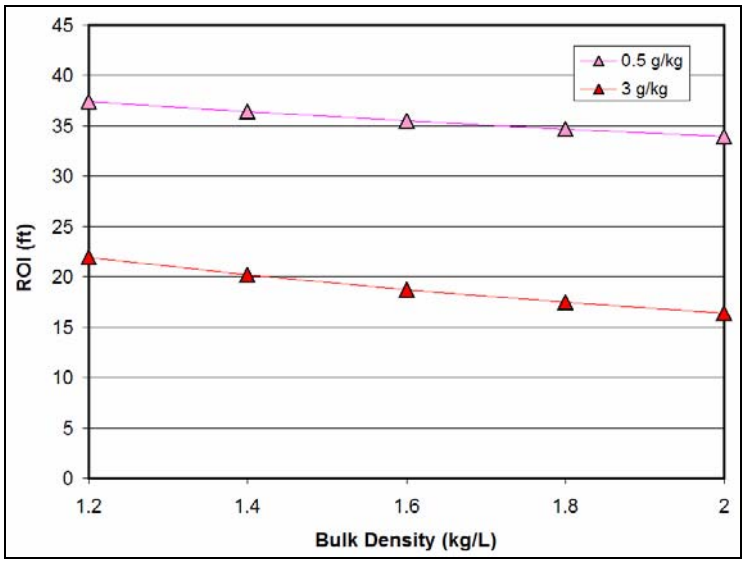

Figure A12-9: ROI Output vs. Bulk Density.

Natural Oxidant Demand (NOD)

The NOD parameter was varied over a wide range, from the conservative case of 0 g/kg NOD to a maximum of 20 g/kg (an unrealistically high value which would generally preclude the use of permanganate treatment). ROI decreases dramatically as NOD values increase from 0 to ~3 g/kg (Figure A12-10). For NOD values greater than ~10 g/kg, though, ROI output remains fairly stable. The lowest possible NOD value that will produce a valid ROI value within the model domain is 0.22 g/kg NOD, but the model domain was extended in order to determine the ROI result for the case of NOD = 0 g/kg. Model sensitivity is greatest for lower NOD values such as 3 g/kg or less (probably the range that most users will be investigating), while NOD inputs of 10 g/kg or greater produce less sensitive output.

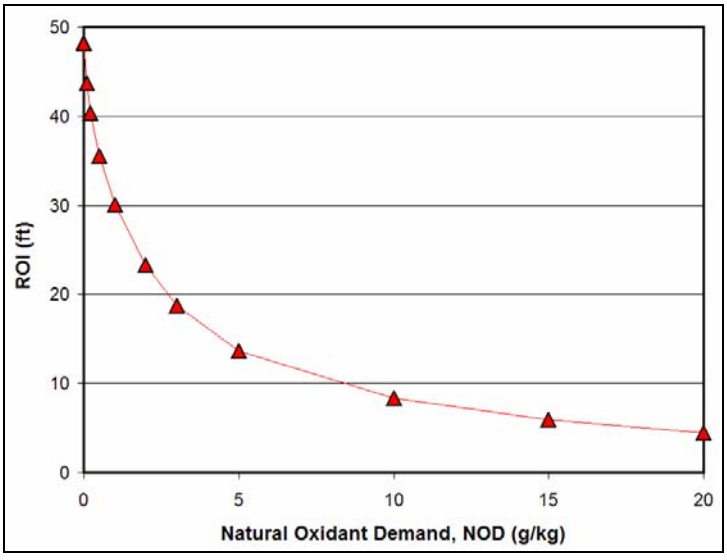

Figure A12-10: ROI Output vs. Natural Oxidant Demand (NOD).

Fraction Instantaneous NOD

This parameter was varied over its entire range, from 0% of the total NOD to 100%. As in the case of the porosity parameter, the ROI response differs depending on the magnitude of the total NOD (Figure A12-11). For the low total NOD (0.5 g/kg) case, ROI correlates inversely with the fraction of instantaneously

reacting NOD (as the fraction increases, ROI decreases). This is expected, since increasing the fraction of instantaneous NOD is similar to increasing the total NOD, which as discussed above decreases the ROI.

For high values of NOD (e.g., greater than 3 g/kg), the ROI is relatively insensitive to fraction of instantaneous NOD. This result is likely due to the presence of adequate reactive NOD, even with small fractions of instantaneous NOD, to drive the complete consumption of the oxidant. For user guidance, it appears that the fraction of instantaneous NOD may play a larger role for low NOD cases (likely to be the majority of modeled scenarios), and without adequate data for this parameter, users may want to investigate the full range of possibility.

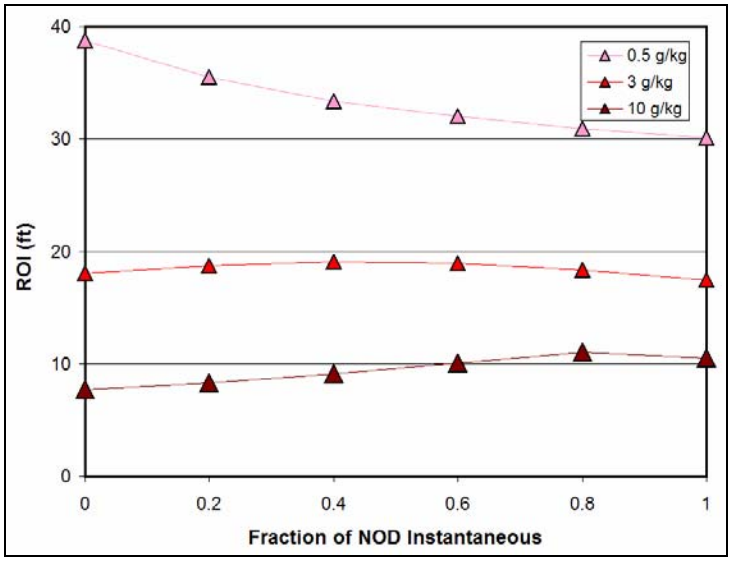

Figure A12-11: ROI Output vs. Fraction Instantaneous NOD.

Slow NOD Reaction Rate

This parameter was varied between values of 1×10^{-4} and 1 L/mmol/day, as these rates bracket the range over which the model is sensitive to changes in this parameter. ROI output decreases as this slow reaction rate increases, reflecting the faster consumption of a larger portion of available NOD (figure A12-12). In reality, the model is only sensitive to this parameter over a two-order of magnitude range of 1×10^{-3} and 1×10^{-1} L/mmol/day. For rate values greater or less than this specified range, the model predicts very little change in ROI. Within these bounds, ROI varies by ~10 feet for both the 0.5 and the 3 g/kg NOD cases.

Initial guidance (1×10^{-4}–1×10^{0} L/mmol/day) for selection of this parameter was based on a survey of the literature. From that information and the results of this brief analysis, users should be directed to choose rate values between 1×10^{-3} and 1×10^{-1} L/mmol/day, unless they have more accurate data to indicate otherwise.

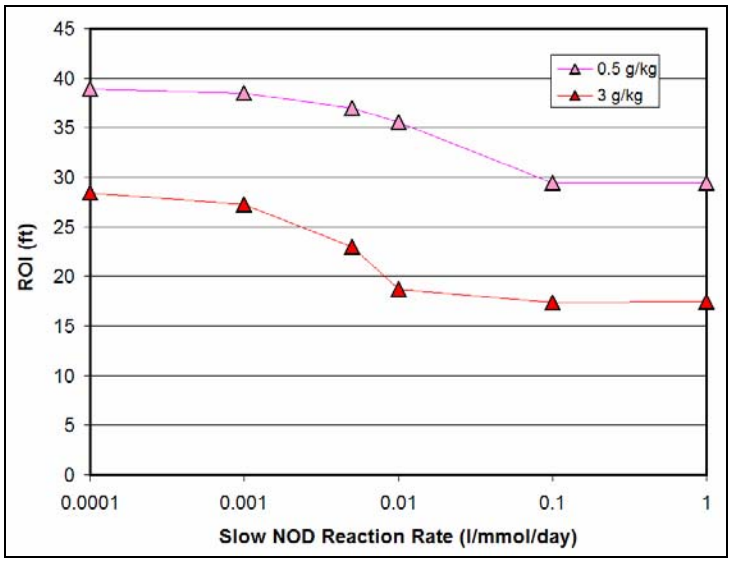

Figure A12-12: ROI Output vs. Slow NOD Reaction Rate (note log scale of x-axis).

Initial Contaminant Concentration

The initial contaminant (TCE) concentration was varied between values close to the MCL level (6×10^{-5} mg/L) and a maximum of 6×10^{3} mg/L, which is greater than the solubility of the contaminant (~1,100 mg/L), indicating the presence of pure-phase contaminant. The model is completely insensitive to initial TCE contaminant levels below ~100 mg/L (Figure A12-13). In fact, for concentrations typical of dissolved-phase site contamination conditions, we can assume that this input parameter has little to no impact on model output. Only when initial TCE concentrations are higher than ~100 mg/L, which as the guidance indicates may be input to mimic pure-phase conditions, does the ROI change in any meaningful amount. This model insensitivity except at extremely high contaminant concentrations underscores the relative importance of NOD and oxidant consumption over contaminant reaction in the model formulation.

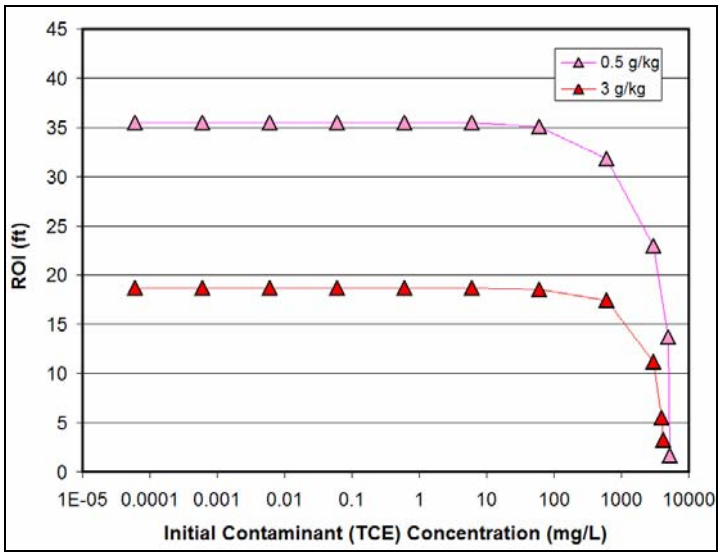

Figure A12-13: ROI Output vs. Initial Contaminant Concentration (note log scale of x-axis).

Injected Oxidant Concentration

The injected oxidant concentration was varied between 1,000 mg/L (or the minimum value necessary to persist within the target ROI for the required number of days) to a maximum of 40,000 mg/L (or that which produces an ROI value within the 40-feet model domain). As one would expect, ROI values increase as the concentration of injected oxidant increases, and this increase is more significant at relatively low oxidant concentrations (~1,000–5,000 mg/L) (Figure A12-14). This wide variation in ROI output is to be expected as this parameter intuitively has a large impact on the model outcome. Additionally, this parameter is one of the main "knobs" that users will most likely adjust during simulations, so they will naturally explore the range of output values.

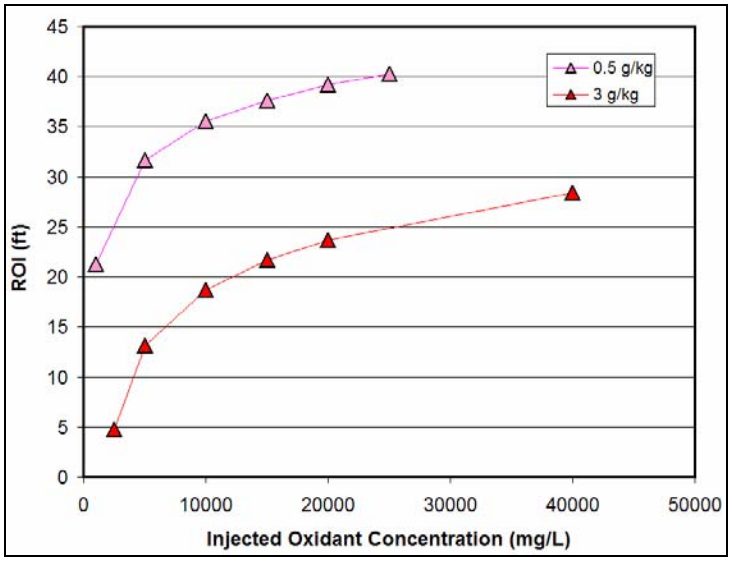

Figure A12-14: ROI Output vs. Injected Oxidant Concentration.

Injection Rate

This parameter was varied over a range of 200 to 15,000 gal/day. As expected, ROI output increases as the injection rate increases. The rate of ROI increase is slightly stronger at lower injection rates (~200–1,000 gal/day) (Figure A12-15).

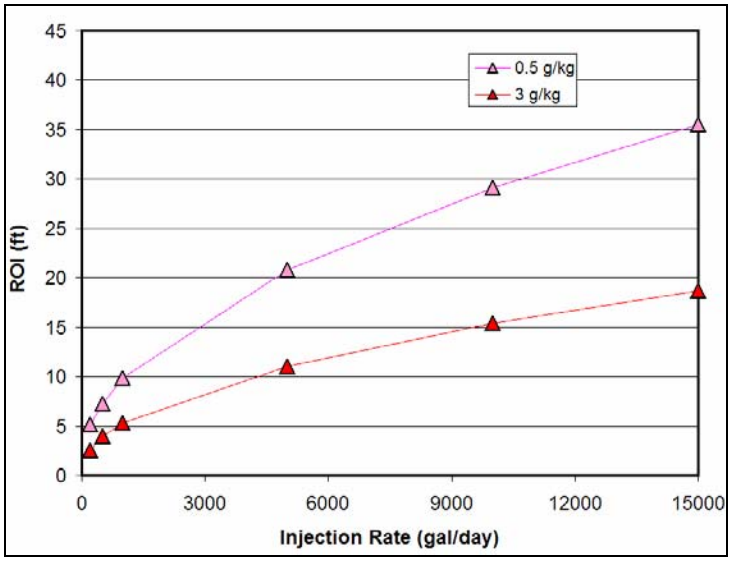

Figure A12-15: ROI Output vs. Injection Rate.