



Appendix 7. Constrained-optimization example

Constrained optimization is an approach to management of water use by using a groundwater flow model. In this context, groundwater pumping rates are adjusted to meet an optimal condition, subject to a set of constraints. For example, the optimal condition may be to maximize pumping subject to the constraint of maintaining a specified flow rate in a stream. Or, if streamflow is already lower than a constraint value, we can minimize the reduction in pumping necessary to achieve streamflow at or above the constraint value.

This approach uses the mathematical technique of linear programming to solve the constrained minimization/maximization problem. Details of the implementation in MODFLOW and of the algorithm are available in Ahlfeld and others (2005) and references within.

Setting a constraint

A potential management scenario was set up to evaluate potential uses of MODFLOW-GWM in the Little Plover River area. This scenario is based on using the steady state model for 2013 and the public rights baseflow value at the Eisenhower gage as a constraint. The steady-state model is conservative with respect to streamflow and is more efficient to use for optimization. A similar approach can be extended to the transient model explicitly, but at a higher computational cost beyond the scope of this proof-of-concept. While multiple constraints can be evaluated at once, Eisenhower was chosen for evaluation both to simplify interpretation to a single constraint and because the Eisenhower gage is the only location on the Little Plover River with continuous streamflow measurements. The public rights

flow value is 4.0 cubic feet per second (cfs) with a corresponding baseflow value of 80 percent or 3.2 cfs. In 2013, the measured baseflow at Eisenhower was 1.95 cfs and the optimal modeled value was 2.09 cfs, below the public rights baseflow value. To achieve flow at or above the public rights baseflow value, one management option to explore is to minimize the reduction of existing pumping required to achieve the public rights baseflow. The deficit in flow is 1.11 cfs or 498 gpm. Based on the depletion-potential calculations described in the Example 3: Depletion-Potential Mapping section of the main report, only wells with depletion potential (relative to the Eisenhower gage) greater than 0.25 were included as managed wells. Wells with depletion potential less than 0.25 obtain 75 percent of the water they extract from surface water features other than the Eisenhower gage on the Little Plover River. As a result, they are not considered in the management with respect to the Little Plover.

Decision variables

To meet the constraint, pumping in wells is adjusted to minimize the reduction in pumping. Managing each well individually is computationally expensive and also results in numerical instability as the response due to a single well may be small. As a result, most optimization efforts group wells into “decision variables,” or groups of wells to manage as a unit. Several approaches to grouping the pumping wells into decision variables were explored, including:

- Create groups based on clustering analysis (discussed below) and apply the same reduction to each well within a group.
- For each clustering scenario, we can either limit the reduction in pumping (for example, no single group or well can reduce its pumping rate more than 35 percent) or we can allow groups to be reduced by 100 percent (in other words, be completely shut off).

Table 1 outlines the names of the scenarios evaluated and shows the total amount of pumping reported for the managed wells in each scenario and the amount of reduction required to meet the public rights baseflow constraint. The following sections further define and present the results of each of these approaches.

Single-decision variable

Applying the same reduction to all wells in the management group is one concept of an equitable solution, spreading the change to water consumption equally among all wells in the area. However, this does not account for the fact that not all wells impact the stream equally. As a result, it is expected that more reduction will be required in aggregate than more flexible approaches. The total pumping without optimization for the 75 managed wells was 4,527 gpm. At this level of pumping, the flow rate at Eisenhower is 1.11 cfs lower than the constraint value of 3.2 cfs. This is a deficit in flow of 498 gpm. The reduction in pumping required (if applied evenly across all wells) to maintain streamflow at the constraint value of 3.2 cfs, however, is 1,356 gpm. Spread evenly across all managed wells, this represents a reduction of 30 percent applied to each well. When apportioning the same reduction among

- Using a single decision variable with all wells in it. This applies the same reduction to each well in the area equally, regardless of depletion potential.

**Table 1.** Percent reduction targets

Scenario	Number of clusters	Max. reduction permitted per group (%)	Baseline base pumping (gpm)	Reduction in pumping (gpm)	Percent reduction (%)
Single-decision variable	1	100	4,527	1,356	30
35% max reduction	20	35	4,527	1,155	26
100% max reduction	20	100	4,527	908	20
35% max reduction, non-irrigation separate	20	35	4,477	1,192	27
100% max reduction, non-irrigation separate	20	100	4,477	883	20

all wells, some wells that are far from the gage with relatively low depletion potential must reduce by the same amount as wells that are much closer with higher depletion potential. Clearly, while the burden of pumping reduction throughout the basin is borne equally in a sense, the overall reduction of pumping in the basin is high. By managing smaller groups and focusing on wells with higher depletion potential, the deficit in flow can be made up with an overall lower reduction in pumping from the whole basin. Both the overall reduction in pumping throughout the basin is expected to be lower, and the burden is focused on fewer wells that have the most impact rather than requiring all wells in the area to have reduced flow. Both approaches can be interpreted as fair by different criteria.

Cluster analysis

Dividing the entire set of wells into management groups can be done in many ways. Of course, issues such as property ownership and crop value should be taken into consideration when creating groups as decision variables. To illustrate the general characteristics of the methods available, however, we looked for a systematic way to make groups and settled on cluster analysis.

K-means clustering (Ahlfeld and others, 2005; Hastie and others, 2009) is a cluster-analysis technique where wells are grouped into k groups (determined by the user) such that the sum of their numerical attributes is minimized. If those numerical attributes are x-coordinate, y-coordinate, and depletion potential, the result is groups of wells that are located near each other and have similar depletion potential values.

Using the depletion potential values calculated for wells that were actively pumped in 2013 and with depletion potential greater than 0.25, a set of clusters was generated with 20 clusters. The number of clusters was chosen as illustrative and other numbers could be implemented easily. Additionally, two variants of this group were generated—one with all wells grouped into clusters, and another with non-irrigation wells managed as individual decision variables and groups made up only of irrigation wells. Figure 1 shows the 20-cluster arrangement created with k-means clustering. Figure 2 shows the 20-cluster arrangements with non-irrigation wells managed individually (the non-irrigation wells are assigned values 20 and greater). For non-irrigation wells, if the reported pumping rate was less than 75 gpm in the steady-state model for 2013, the

well was left out of the management scenario to limit the number of decision variables. Such a limitation was not used for the other scenarios as the wells with lower pumping rates were incorporated into groups.

In addition to grouping wells into decision variables, a decision must be made regarding how much pumping can be reduced in each decision variable to meet the constraint. We evaluated two cases: one where the maximum reduction was 35 percent, and another where the maximum reduction was 100 percent.

Table 1 outlines the scenarios showing the various combinations of the number of groups (including cases where non-irrigation wells were managed separately) and the maximum pumping reduction allowed. In the following sections, the results of these scenarios are presented and discussed.

20 clusters

In these two scenarios, k-means clustering was used to define 20 groups for management and the maximum pumping reduction allowed in each group was set at either 35 percent and 100 percent. In other words, in the 35% Max Reduction scenario, no pumping reduction for a group was allowed to drop by more than 35




Figure 1. Arrangement of wells into 20 clusters.




Figure 2. Arrangement of wells into 20 clusters with non-irrigation wells kept separate.



percent. In the 100% Max Reduction scenario, groups were allowed to be completely turned off, thus simulating a more binary response with wells either on or off. Figures 3 and 4 show the amount and locations of pumping that must be reduced to maintain the constraint of the public rights baseflow at the Eisenhower gage. Table 1 highlights the overall percent reduction in pumping required over the entire set of managed wells. Relative to the first scenario with all wells subject to the same level of reduction (single-decision variable), the subdivided (grouped) scenarios are more efficient, requiring about 5–10 percent less reduction to meet the constraint. The 35% Max Reduction scenario requires an overall reduction in pumping of 1,155 gpm while the 100% Max Reduction sce-

nario required an overall reduction of 908 gpm. This efficiency is realized by focusing the reduction on wells with the highest depletion potential. Wells with lower depletion potential are obtaining water from sources other than the Eisenhower gage, so reducing their pumping is a less-efficient way to increase flow at the gage. Indeed, a similar difference is related to the maximum amount of reduced pumping allowed. For the 100% Max Reduction scenario, the constraint can be met by turning off wells that are closest to the gage (highest depletion potential). However, if the reduction is limited to 35 percent, the constraint can only be met by also reducing pumping in wells with lower depletion potential in addition to wells with highest depletion potential, thus the overall reduction in pumping is higher (25.5 percent vs. 20.1 percent).

20 clusters, nonirrigation wells separate

These two scenarios are the same as the previous two with the exception that non-irrigation wells were not considered as part of the k-means clustering process and individual non-irrigation wells with a pumping rate less than 75 gpm were not managed. Table 1 shows that overall reduction in pumping required is similar to the 20 cluster scenarios. Figures 5 and 6 graphically show the amount and locations of pumping that must be reduced to maintain the constraint at the Eisenhower gage with maximum allowable reduction at a single well of 35 percent and 100 percent, respectively. The results with non-irrigation wells managed separately are similar to the previous cases—particularly for the maximum reduction of 35 percent. The 35% Max Reduction Non-Irrigation




Figure 3. Optimal rates of pumping reduction for 20 clusters, constrained to meeting the public rights baseflow at Eisenhower, when maximum allowable flow reduction is 35 percent.




Figure 4. Optimal rates of pumping reduction for 20 clusters, constrained to meeting the public rights baseflow at Eisenhower, when maximum allowable flow reduction is 100 percent.



Separate scenario and the 100% Max Reduction variant require total reduction in pumping of 1,192 gpm and 883 gpm, respectively. The patterns observed in figures 3 and 5 and are similar, although in figure 5, where non-irrigation wells are separated out, greater reduction is achieved in some individual wells, which shifts the pattern of pumping reduction around. In the 100% Max Reduction case, the patterns are more distinct because the flexibility of more management groups (because some non-irrigation wells are pulled out from groups south of the river) allows for focused reductions. The main difference here is that urban or industrial wells are not expected to be managed by the same entities as irrigation wells. Separating out the groups allows decisions variables to be composed of wells that are more likely managed similarly with one another and by consistent managers.

Conclusions

Constrained optimization is a powerful tool to objectively adjust well pumping to meet a constraint of streamflow. Wells are managed such that the most impactful wells—those with the highest depletion potential—are targeted for reduction so the streamflow constraint is maintained with the least amount of reduction possible. The arrangement of decision variables, either grouping all wells into one or splitting into smaller groups using cluster analysis or other techniques, has profound ramifications both on how much each well is affected, and how much total reduction in pumping is required.

In managing a basin, stakeholders and decision makers must determine a concept of fairness to apply. Requiring each well to reduce by the same percent spreads the burden throughout the managed wells, but

results in many more wells being reduced and a less efficient allocation of that reduction than when smaller groups are managed. However, at the other extreme, if individual wells or small groups of them are allowed to reduce by 100 percent, some wells will shoulder all of the burden while others will not be affected at all. As a basin, mechanisms to compensate the most impacted users would need to be developed which is beyond the scope of this work.

References

- Ahlfeld, D. P., Barlow, P. M., and Mulligan, A.E., 2005, GWM—A groundwater management process for the U.S. Geological Survey modular groundwater model (MODFLOW-2000): U.S. Geological Survey Open-File Report 2005-1072, 124 p.
- Hastie, T., Tibshirani, R., and Friedman, J.H., 2009, The elements of statistical learning: Data mining, inference, and prediction: New York, Springer, Springer Series in Statistics, 745 p.




Figure 5. Optimal rates of pumping reduction for 20 clusters, constrained to meeting the public rights baseflow at Eisenhower, when maximum allowable flow reduction is 35 percent. Non-irrigation wells managed individually.




Figure 6. Optimal rates of pumping reduction for 20 clusters, constrained to meeting the public rights baseflow at Eisenhower, when maximum allowable flow reduction is 100 percent. Non-irrigation wells managed individually.