A Simulated Annealing Approach to Irrigation

February 6, 2006

Abstract

A significant percentage of the Earth's surface is farmland. Unfortunately, nature does not usually supply enough rainfall for the types of plants that most farmers want to grow. For this reason, there has been a fair amount of engineering research devoted to irrigation technology. Much of this work, however, appears to be based on empirical data and crude estimates. One particular instance of this is the design and placement of sprinkler systems on fields. We are interested in developing a precise mathematical model/algorithm to determine the design and schedule of a hand-move irrigation system for small fields. The goal is to minimize the amount of time required by the farmer to move the sprinklers while ensuring that no portion of the field receives either too much or too little water.

The flow of water from a pump through a pipe system to a set of sprinklers is governed by Bernoulli's equation and the equation of continuity. To model the distribution of water out of each sprinkler we calculate the trajectory of water droplets emitted from the nozzle as they split into smaller drops and are affected by air resistance. The precipitation rate distribution from a pipe-set is then determined by superposition of the individual distributions.

Based on the parameters determined by these models, we use a simulated annealing algorithm to find the configuration of pipe-set locations that produces the most uniform precipitation distribution across the field for a variety of parameters. An instance of the traveling salesman problem is then needed to determine the optimal order to move the pipe-set through the specified locations. For low parameters we are able to solve this problem with an exhaustive search, but for larger values we employ a heuristic approach known as the Christofides algorithm. We estimate the walking time required by the farmer to move the pipe-set in the order produced from this algorithm. For each choice of parameters (number of sprinklers and number of pipe-set moves), we rule out the configuration if flooding or drought occurs. Finally, we recommend to the farmer the configuration that has minimal upkeep cost (in terms of walking time) that has not been ruled out. This configuration is a slightly skew square patter in which a pipe-set consisting of 3 sprinklers is placed horizontally and moved around the corners of the field after running in each position for 11.25 hours. This configuration has a maximum precipitation rate of 0.26cm/hr and a minimum accumulation of 2.99cm in four days.

Contents

Figure 1: The pipe-set of a typical hand-move irrigation system (photo obtained from USDA website[1])

1 Introduction

A common device for irrigating a field is a *hand-move* irrigation system. This is a set of straight aluminum pipes, each with a sprinkler extending vertically into the air, that can be joined together to produce a line of sprinklers known as a *pipe-set* (see Figure 1). The pipe-set is placed in the field and at regular intervals the farmer disassembles the pipes, moves them one at a time to a new location, and then reassembles them. This process of moving the pipe-set is time-consuming, so the farmer is inclined to perform as few moves as possible; however, if the pipe-set remains in any one place too long the crops in that area will become hyper-saturated whereas those in other parts of the farm will suffer from dehydration.

1.1 The Problem

Our goal is to produce a model of this type of irrigation system and develop an algorithm based on this model to determine the optimal irrigation schedule for a given field. In particular, we wish to find the location and period of pipe-set placements in a field of size 80m×30m that minimizes the time required by the farmer subject to the following constraints:

Constraints:

- No part of the field is to be over-watered (receive more than $0.75\frac{cm}{hr}$ on average) or under-watered (receive less than $2cm$ in 4 days).
- The total length of the pipe-set is 20m, the inner diameter of each pipe is 10cm, and the nozzle diameters are 0.6cm.
- The volumetric flow rate of the water pump may not exceed 150 liters per minute, and a constant pressure of 420kPa is maintained at the inlet of the pipe-set.

As part of this problem we also determine the optimal number of sprinklers to place on the pipe-set. In addition to these constraints, we make a few basic assumptions about the setup:

Assumptions:

- Wind patterns are rather unpredictable and thus do not affect the choice of irrigation scheme.
- The orientation of the pipe-set is fixed throughout the schedule to minimize the difficulty of movements for the farmer.
- The sprinklers are spaced uniformly across the pipe-set.
- The farmer moves the pipe-set at *regular* time intervals, so that the sprinklers are running in each location for the same amount of time.

1.2 Our Approach

The general problem consists of two components: the model and the algorithm. We break the model up into a few submodels. We also divide the optimization algorithm into several phases to reduce the number of parameters under consideration at each stage of computation. This approach is outlined as follows:

Model:

1. Using basic fluid mechanics, determine the velocity and volume of water at the sprinkler nozzles induced by the pump.
2. Calculate the distribution of water emitted from a single sprinkler, incorporating air resistance and variation in water droplet size/velocity/angle.
3. Produce the total precipitation rate distribution of the pipe-set by superposition of multiple sprinklers.
4. Estimate the time required for the farmer to move a pipe-set from one site to another.

Algorithm:

1. For a fixed number of sprinklers and pipe-set sites, call them s and l respectively, use a simulated annealing optimization process to determine the configuration of pipe-sets that produces the most uniform precipitation rate distribution throughout the field.

2. Create a table of the configurations resulting from step 1 as s and l vary over a specified range.

3. For each such configuration find the optimal ordering of pipe-set sites to minimize the farmer's walk time (i.e. solve a particular instance of the Traveling Salesman Problem) and create a table of these cost times.

4. Determine which configurations in the table from step 2 satisfy the minimum and maximum water requirements (i.e. rule out schedules in which flooding or drought occur).

5. Of the configurations not ruled out in step 4, choose the configuration with the minimal total walk time calculated in step 3.

2 Water Pump and Pipes

The two principles needed to understand the behavior of fluid in a pipe (see [2]) are *conservation of mass*, which is expressed in the equation of continuity:

$$vA = constant \tag{1}$$

where v is the velocity and A is the cross-sectional area at a given location in the pipe, and *conservation of energy*, which is essentially Bernoulli's Equation:

$$p + \frac{1}{2}\rho v^2 = constant \tag{2}$$

Here p is pressure and ρ is the density of water. Technically there is also a term $\rho g y$ in Bernoulli's equation representing potential energy due to gravity, but everything in our problem occurs at the same altitude y, so this term drops out.

2.1 A Pipe with One Sprinkler

Consider the pipe in Figure 2 with cross-sectional area A_1 attached to a pump at one end and a sprinkler with nozzle area A_2 at the other. Let v_1 be the velocity of the water flowing near the pump and v_2 the velocity at the nozzle, and let p_1, p_2 be the pressures at the corresponding locations. Then by Bernoulli's equation we have that

$$p_1 + \frac{1}{2}\rho v_1^2 = p_2 + \frac{1}{2}\rho v_2^2. \tag{3}$$

Using the equation of continuity, we get

$$v_1 = \frac{v_2 A_2}{A_1}, \tag{4}$$

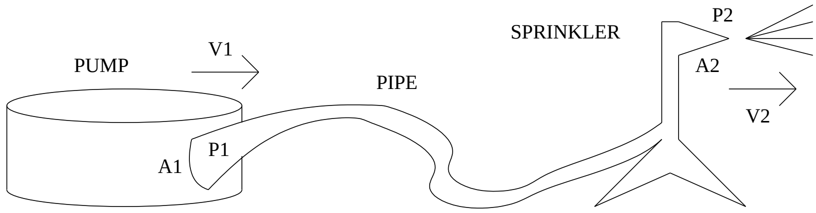

Figure 2: Water flowing from a pump to a sprinkler

so we can solve for the velocity at the nozzle:

$$v_2 = \sqrt{\frac{p_1 - p_2}{\frac{1}{2}\rho(1 - (\frac{A_1}{A_0})^2)}}. \tag{5}$$

Now p_1 is the given pressure at the pump (in our case 420kPa) and p_2 is atmospheric pressure (101.325kPa), so this formula allows us to compute an actual numerical value for the nozzle velocity. There are, however, two potential limitations to this equation:

1. **Pump capacity**

   By the equation of continuity, the volumetric flow rate $q = vA$ remains constant throughout the system. Since the pump has a maximum flow rate q_{max} (in our case $q_{max} = 150\frac{liter}{min}$), the flow rate through the nozzle $q_{noz} = v_2 A_2$ must satisfy $q_{noz} \leq q_{max}$. Therefore, if the nozzle velocity given in equation (5) exceeds $\frac{q_{max}}{A_2}$ it must be replaced by $\frac{q_{max}}{A_2}$.

2. **Pressure loss due to friction**

   In the real world, a certain amount of water pressure is lost due to friction in the pipe. This has been found to behave according to the Hazen-Williams formula[3]:

   $$\Delta p = -4.55L\frac{(q/c)^{1.852}}{d^{4.87}} \tag{6}$$

   where L is the length of the pipe, d is the diameter, and c is a roughness coefficient ($c = 120$ for aluminum pipes such as the ones in our pipe-set). Fortunately, the loss in a 20m pipe-set according to this equation is 0.334kPa, which is insignificant relative to the other pressures involved.

2.2 Multiple Sprinklers

Now suppose that there are n sprinklers placed uniformly along a pipe. Because friction is negligible, the flow rate through each nozzle is distributed uniformly among the sprinklers. Moreover, the pressure drop across each sprinkler induces a current independent of the others, so the total flow through all the sprinklers is simply the sum of those predicted by the previous case – as long as the pump capacity is not exceeded. Putting all this together shows that,

the velocity at each nozzle of an n-sprinkler pipe-set irrigation system is:

$$v_{noz} = \begin{cases} \sqrt{\frac{p_{pump} - p_{atm}}{\frac{1}{2}\rho(1 - (\frac{A_{noz}}{A_{pipe}})^2)}}, \text{ if } nq_{noz} \le q_{max} \\ \frac{q_{max}}{nA_2}, \text{ otherwise} \end{cases} \tag{7}$$

3 Flow Distribution from the Sprinklers

We now determine the precipitation distribution from the pipe-set.

3.1 Modeling Droplets

We model the spray of water as a collection of individual, spherical droplets, each considered using the model developed in [5]. In particular, we assume that the motion of a droplet is given by

$$\frac{d\vec{v}}{dt} = -g\hat{z} - \frac{\eta |v|^2 \hat{v} r^2}{m} \tag{8}$$

where $g = 9.8\frac{m}{s}$ is the acceleration due to gravity, $\vec{v}$ is the velocity vector of the droplet, m is its mass, r its radius, $\eta = 0.855\frac{kg}{m^3}$ a constant of proportionality, and $\hat{v} = (v_1, v_2)/\sqrt{v_1^2 + v_2^2}$ the unit vector in the direction of the velocity (here v_1 and v_2 denote the horizontal and vertical components of the velocity, respectively). Splitting this vector equation into components and using difference equations, we can rewrite this equation as

$$dv_1 = -\frac{\eta r^2 v_1 |v|}{m} dt \tag{9}$$

$$dv_2 = \left(-g - \frac{\eta r^2 v_2 |v|}{m}\right) dt \tag{10}$$

We then apply Euler's Method to these equations, with $dt = 0.001$, to determine the velocity of the droplet. Applying Euler's method again to the equation

$$d\vec{x} = \vec{v} dt \tag{11}$$

gives the path of the droplet.

Next, we assume that droplets break up at a rate $\lambda = \lambda_0 |v| r^2$, where $\lambda_0 = 5000\frac{\text{breakups}}{\text{sec m}^3}$ is a fixed constant. When a droplet A splits, the mass is divided among two new droplets, B and C, with the mass of the new droplets determined from a uniform distribution. Each new droplet is given a slightly different vertical velocity component from that of its parent. As smaller droplets have a lower terminal velocity, a droplet that splits will fall closer to the sprinkler than one that does not.

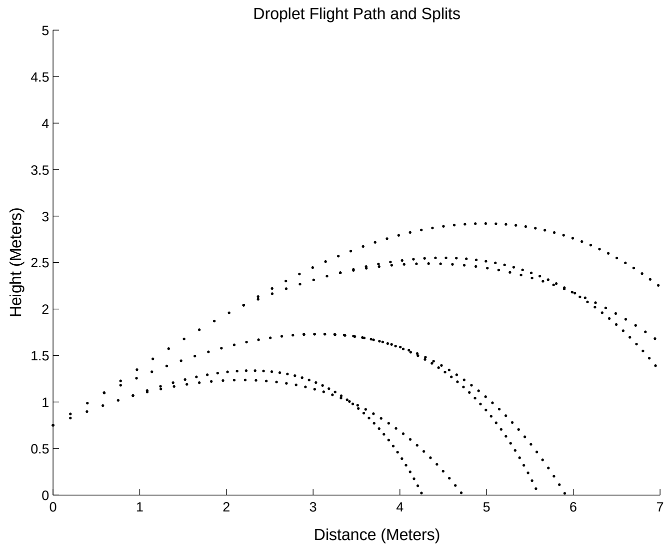

Figure 3: The flight-path of a single droplet as it splits into several drops.

3.2 Modeling a Sprinkler

We model the flow of water from the sprinkler as a collection of droplets. Unfortunately, little data is available regarding the characteristics of the water distribution. For the sake of definiteness, we assume that:

- The sprinkler is 0.45 meters above the ground (estimated from Figure 1).
- The sprinkler launches droplets with velocity given by a normal distribution with mean velocity v determined by equation (5) and standard deviation $0.4v$.
- The droplets have an initial direction uniformly distributed between 0 and 30 degrees.

The allows us to estimate the water flow from a single sprinkler (see Figure 4).

Because the sprinkler has a *rotating* spray nozzle, the orientation of the droplet trajectory is slowly rotated around the sprinkler. We assume that the angular velocity of the nozzle is small enough that it does not affect the individual droplet motion. This gives a distribution of droplet landing points centered at the sprinkler as depicted in Figure 5, and we can use this to determine the precipitation flow rate on the farm near this sprinkler (Figures 6 and 7). For low water velocities ($10\frac{m}{s}$ or less), the distribution of water is approximately bell-shaped. For higher velocities, the distribution becomes more volcano shaped.

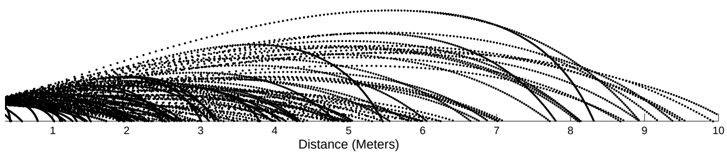

Figure 4: Flow of water from a sprinkler.

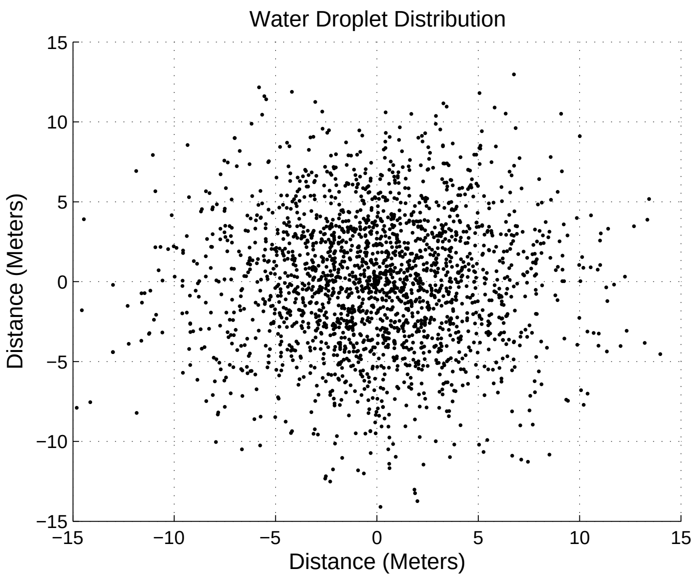

Figure 5: Landing positions of droplets relative to the sprinkler.

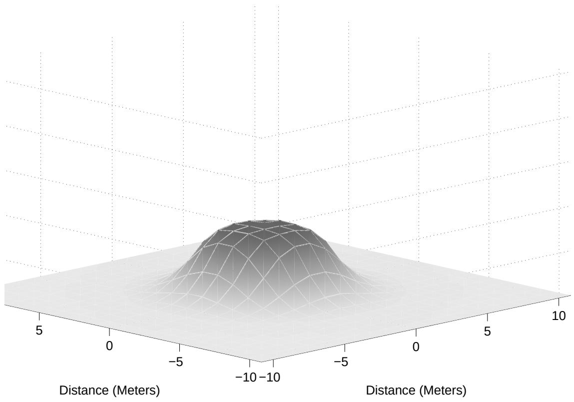

Figure 6: Distribution of water relative to the sprinkler, v = 6 m/s.

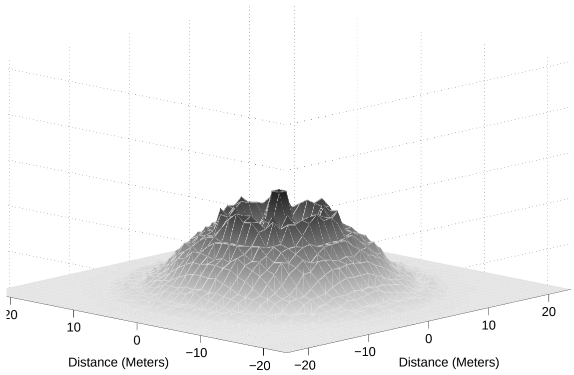

Figure 7: Distribution of water relative to sprinkler, v = 12 m/s.

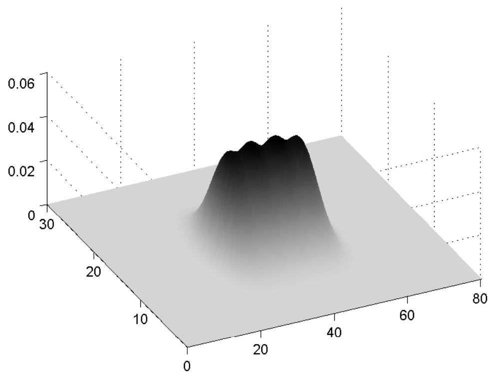

Figure 8: Superposition of four sprinklers to produce the distribution from a pipe-set.

3.3 Flow from a Pipe-set

We can now produce the precipitation rate distribution from the entire pipe-set simply by super-positioning distributions of the individual sprinklers calculated in the previous section according to equation (5). The result is pictured in Figure 8.

3.4 Farmer Movement Times

The time required for the farmer to move the irrigation system depends on various factors. First, the farther he has to move the system, the longer it will take. Second, we assume that each sprinkler is connected to its own pipe, with the sum of the lengths of the pipes being fixed at 20 meters. Thus a large number of sprinklers in the irrigation system results in significant assembly and disassembly times and repeated trips per movement of the irrigation system. However, a small number of sprinklers results in long, unwieldy pipe assemblies. In particular, the time required to move the irrigation system is the time required to disassemble and reassemble, plus the time required to carry each sprinkler to the next site, plus the time required to walk from the new site back to the old one. To model the time required of the farmer to move the irrigation system, we used the following equation:

$$T = 2nt + \beta n d e^{\alpha l} + \frac{nd}{v} \tag{12}$$

where T is the total time required for the farmer to move the irrigation system, n is the number of sprinklers, t is the time necessary to assemble one pipe/sprinkler system, l = length of pipes = $20/n$, $v = 1$ m/s is the natural walking speed of the farmer, and α and β are the fixed constants. For our model we took $\alpha = 1/5$

and $\beta = 1/10$, as these values gave a reasonable movement time of about 30 minutes.

4 Irrigation Scheduling Algorithm

We would like to find the number of sprinklers s on a pipe-set and an ordered list of coordinates $C = ((a_1, b_1), \ldots, (a_l, b_l))$ representing sequential pipe-set sites so that the time required of the farmer to adhere to the irrigation schedule is minimal, subject to the flood/drought constraint. To accomplish this, we first find the irrigation schedules that yield the most uniform possible precipitation rate distribution throughout the farm for each fixed value of s and $l = |C|$, since these configurations will have the greatest chance of meeting the watering constraints.

Suppose the farmer has decided to use an irrigation schedule with l different pipe-set locations. Since we are assuming that the pipe-set remains at each location for the same amount of time, the average precipitation rate at a given location in the farm is equal to the sum of the precipitation rates due to the pipe-set at each location divided by the number of locations. Formally, if we let $w(x, y)$ be the average precipitation rate at coordinates (x, y) in the field (i.e. the number of inches of water applied each hour averaged over the entire irrigation schedule), then

$$w(x, y) = \frac{w_1(x, y) + w_2(x, y) + \cdots + w_l(x, y)}{l} \tag{13}$$

where $w_i(x, y)$ is the precipitation rate that would be experienced if the pipe-set were fixed at location (a_i, b_i). This allows us to visualize an irrigation schedule simply by super-positioning the precipitation distributions due to each fixed location and then scaling the net result.

4.1 Measuring Uniformity

Throughout this section assume the number of sprinklers s and pipe-set locations l are fixed. Our goal is to choose the locations in a way a that maximizes the uniformity of the average precipitation rate $w(x, y)$. A natural measure of uniformity is given by the squared distance from the mean:

$$\sum_{i,j=1}^{n} (w(x_i, y_j) - \mu)^2 \tag{14}$$

where the field has been discretized into an $n \times n$ grid and $\mu = \frac{\sum_{i,j=1}^{n} w(x_i, y_j)}{n^2}$.

Let $Q((a_1, b_1), \ldots, (a_l, b_l))$ be this squared distance interpreted as a cost function dependent on the coordinates $C = ((a_1, b_1), \ldots, (a_l, b_l))$ of the pipe-set positions. Then we want to find the (a_i, b_i) that minimize the cost Q. To do this, we use a modified Monte Carlo method known as *simulated annealing.*

4.2 Simulated Annealing

The parameter space we are dealing with in this optimization process is too vast to search exhaustively, so we must be content to find an approximate solution. The difficulty is that when minimizing the cost Q numerically, it is easy to get stuck in a local minimum that is far inferior to the global minimum. There is a general method for bypassing this trap known as a *Metropolis* algorithm[4], which we can apply to our case as follows:

1. Place the l coordinates (a_i, b_i) for the pipe-set positions randomly in the field.

2. Compute the cost Q of the current configuration.

3. Randomly perturb all the coordinates by a small amount.

4. Compare the new cost Q' with the old one.

5. If the change in cost $\Delta Q = Q' - Q$ is negative then accept this perturbation and go back to step 2.

   If ΔQ is positive, indicating a step in the wrong direction, then with a certain probability accept the perturbation and return to step 2, otherwise return the points to the previous coordinates and go to step 2.

The probability in the last step depends on the size of ΔQ (the worse the perturbation, the less likely it will be accepted). Simulated annealing gives a more intelligent way to determine which "bad" perturbations should be accepted. The idea is to introduce a temperature parameter $T > 0$. The higher the temperature, the more willing we are to take bad steps. After a specified number of iterations through the loop, the temperature is decreased and the loop resumes. This process continues until the temperature reaches zero. This allows the optimizer to explore the overall landscape at first, and then settle into a specific optimum (hopefully the global one) at the end. Concretely, at each iteration of the loop you should accept the perturbation if

$$e^{\frac{-\Delta Q}{T}} > R(0,1) \tag{15}$$

where $R(0,1)$ is a random number chosen uniformly from the unit interval. The results of this algorithm for a few different values of s and l are pictured in the appendix (Figures 12-19).

4.3 The Traveling Farmer Problem

Once a configuration of pipe-set locations has been determined, it still remains to decide what order to move the pipe-set through these sites. The goal is to minimize the distance the farmer has to walk during the irrigation schedule. This is an instance of the well-known Traveling Salesman Problem. For most of the configurations we consider, the number l of locations is low enough that an exhaustive search is suitable. However, for $l > 10$ the computational requirements are beyond our resources – and for farms larger than the one we are considering (so that an even greater number of sites is required), an exhaustive search is completely unreasonable. Therefore we suggest a heuristic approach to find a near-optimal solution.

Since the farmer's walking path is measured with the ordinary Euclidean distance, so that the triangle inequality is satisfied, the problem reduces to an easier case known as the *Euclidean TSP*. With this simplification, there is a constant-factor approximation algorithm (due to Christofides[6]) that always finds a tour (path through all the sites) of length at most 1.5 times the optimal tour (see [7]):

1. Find a minimum spanning tree T.
2. Find a minimum distance perfect matching M among those vertices with *odd* degree in T.
3. Form the multigraph G by combining the edges of T and M.
4. Find an Eulerian cycle w in G and, by bypassing the vertices that have already been visited, form a tour t.

It is possible to find an Eulerian cycle in G because there will be an even number of vertices with odd degree in T due to the handshaking lemma, so that all vertices of G have even degree. The intuitive idea behind the correctness of the algorithm is as follows: the length of a tour is an upper bound for the length of a minimum spanning tree because given any tour, the deletion of any edge yields a spanning tree. The complexity of the algorithm is $O(n^3)$, where n is the number of vertices, as opposed to $O(\frac{n!}{2})$ for the exhaustive search.

4.4 Flooding and Drought

For a given configuration, to determine if any portion of the field receives too much water we simply find the location in the field with maximal precipitation rate and see if it is less than $0.75\frac{cm}{hr}$. This value represents the precipitation rate when averaged over the entire scheduling scheme. To ensure that no plants become dehydrated, we need to measure how much water accumulates at the location of minimal precipitation rate in the configuration. For this, we assume that the farmer is only able to keep the sprinklers running for 12 hours per day, since he has other chores to attend to and most likely does not want to run the irrigation system overnight. With this assumption, the 4-day precipitation total is given by w(48-t), where w is the precipitation rate at the point in question and t is the time required to perform all l moves in the schedule (this latter quantity must be subtracted from the total precipitation since the sprinklers are turned off while the pipe-set is being moved). As long as this value is greater than 2cm at all points in the field, we are content that the plants will have sufficient water.

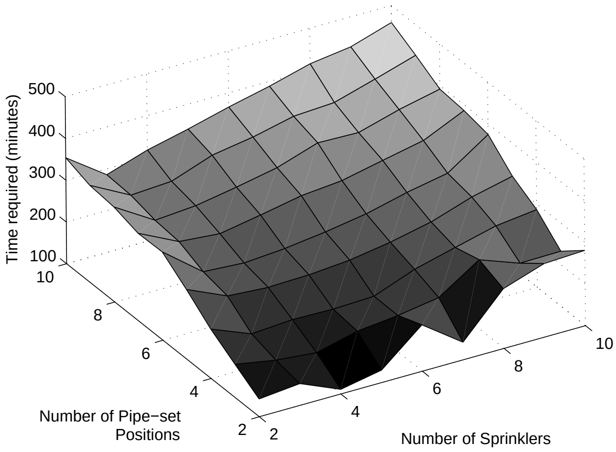

Figure 9: The time required to maintain a schedule for a 96 hour period.

5 Conclusion

Using the data calculated from the water-flow models, we ran the pipe-set placement algorithm for values of s and l both ranging from 1 to 10 to produce a table of configurations that are optimal with respect to precipitation rate uniformity. We then determined an optimal tour for each configuration and estimated the time required by the farmer to maintain the schedule for each 96 hour period. The result of these cost calculations is depicted in Figure 9.

Next, for each configuration we calculated the maximum precipitation rate averaged over the schedule and the minimum precipitation accumulated in a 96 hour period. If either of these values were outside the constraints specified earlier, we eliminated the configuration. The results of this are illustrated in Table 1 (an 'X' indicates a flood/drought situation and an 'O' indicates a valid configuration). The configuration with minimal cost that was not eliminated in this fashion is the one with 4 positions and 3 sprinklers. This configuration is pictured in figure 10. We now summarize the salient features of this configuration:

- Maximum precipitation rate: $0.26\frac{cm}{hr}$.
- Minimum accumulation over 96 hours: $2.99cm$.
- Total time required by the farmer to move sprinklers in a 4-day period: 3 hours.

The instructions for the farmer to use this schedule as follows:

Table 1: Valid configurations, as a function of number of pipe-set positions (vertical) and number of sprinklers (horizontal).

	1	2	3	4	5	6	7	8	9	10
1	X	X	X	X	X	X	X	X	X	X
2	X	X	X	X	X	X	X	X	X	X
3	X	X	X	X	X	X	X	X	X	X
4	X	O	O	X	X	X	X	X	X	X
5	X	X	O	O	X	X	X	X	X	X
6	X	O	O	O	O	O	O	O	X	X
7	X	O	O	O	O	O	O	X	X	X
8	X	O	O	O	O	O	O	O	X	X
9	X	O	O	O	O	O	O	O	O	O
10	X	O	O	O	O	O	O	O	O	X

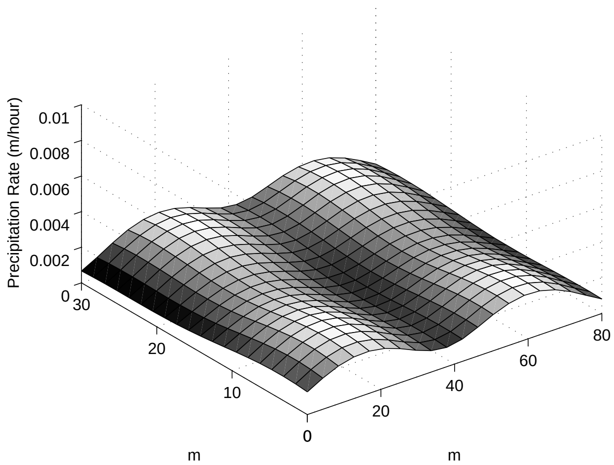

Figure 10: The precipitation distribution for the configuration we recommend.

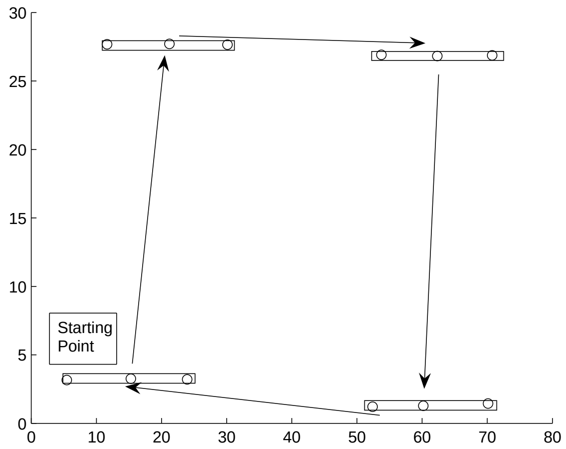

Figure 11: Instructions for the irrigation schedule we recommend.

1. Place the pipe-set horizontally at coordinate (15.3,3.3), where the South-West corner of the field is taken to be the origin.
2. Leave the sprinkler running for 11.25 hours.
3. Move the pipe-set North to (21.2,27.7) and leave it for 11.25 hours.
4. Move the pipe-set East to (62.3,26.8) and leave it for 11.25 hours.
5. Move the pipe-set South to (60.2,1.2) and leave it for 11.25 hours.
6. Move the pipe-set West and repeat this cycle every 96 hours.

6 Appendix

Here we display some of the interesting configurations that were produced by the simulated annealing algorithm.

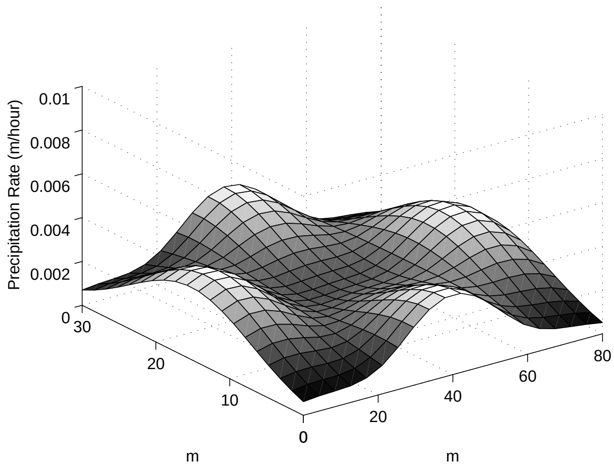

Figure 12: Precipitation rate distribution for 4 sprinklers, 4 positions.

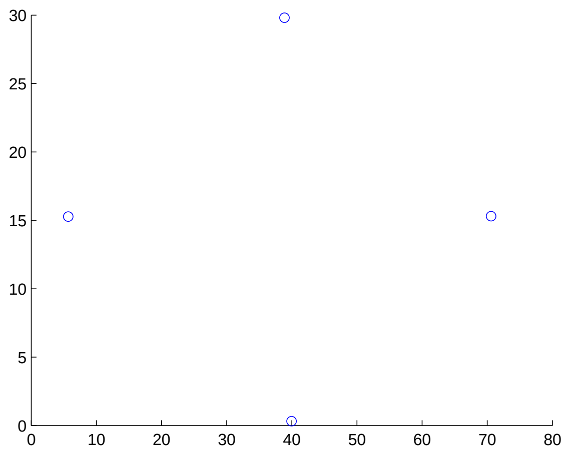

Figure 13: Locations of the pipe-sets.

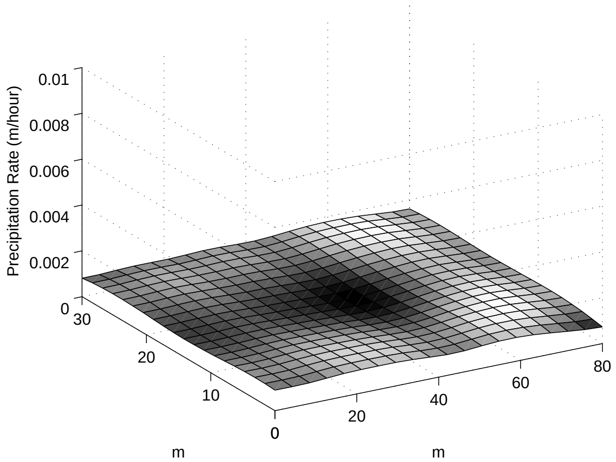

Figure 14: Precipitation rate distribution for 2 sprinklers, 6 positions.

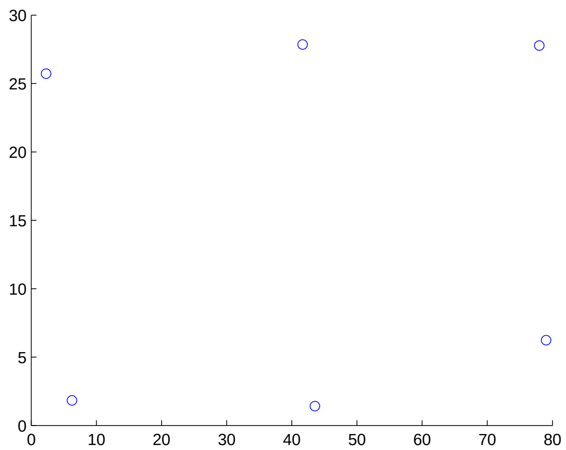

Figure 15: Locations of the pipe-sets.

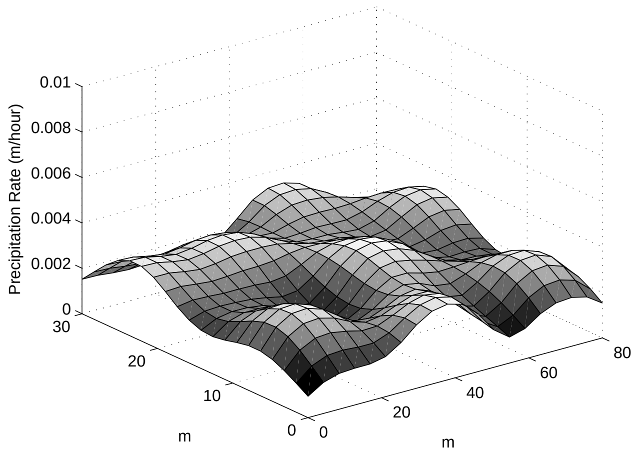

Figure 16: Precipitation rate distribution for 6 sprinklers, 8 positions.

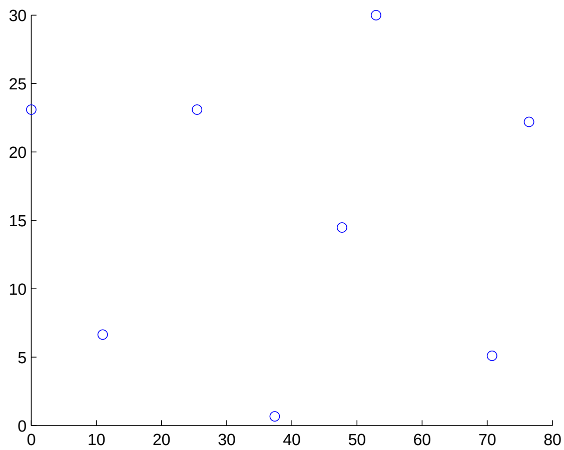

Figure 17: Locations of the pipe-sets.

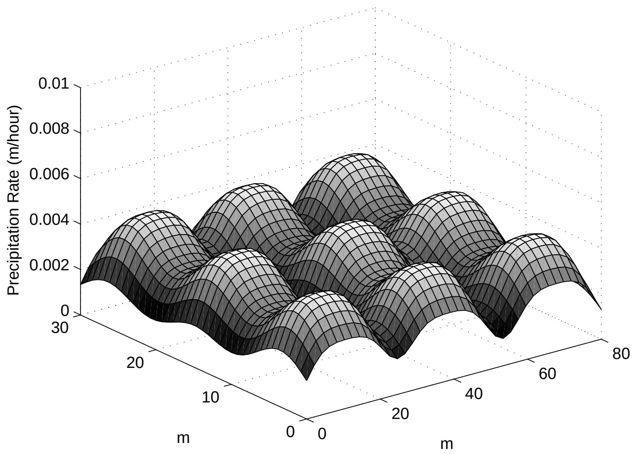

Figure 18: Precipitation rate distribution for 8 sprinklers, 9 positions.

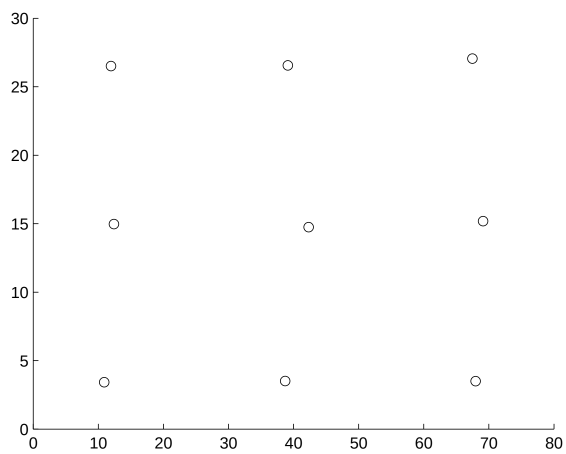

Figure 19: Locations of the pipe-sets.

References

[1] United States Department of Agriculture: www.ers.usda.gov/Briefing/WaterUse/glossary.html. Accessed 2/5/06.

[2] Resnick, Halliday, and Krane. *Physics*. John Wiley and Sons, 1992.

[3] Rain Bird Commercial Irrigation: www.rainbird.com/pdf/ag/FormulasUS.pdf. Accessed 2/6/06.

[4] Otten and van Ginneken. *The Annealing Algorithm*. Boston, MA: Kluwer, 1989.

[5] Card, Esser, Giansiracusa. ''A Foul Weather Fountain,'' *The UMAP Journal* 23 (3) (Fall 2002).

[6] Christofides, ''Worst-case analysis of a new heuristic for the travelling salesman problem,'' Report 388, Graduate School of Industrial Administration, Carnegie Mellon University, 1976.

[7] National Institute of Standards and Technology: http://www.gist.gov/. Accessed 2/6/2006.