Project 4: Olympic Road Race

Group 3

Sasha Davidov, Becky Plummer, & Anya Robertson

13 December 2004

Introduction

The Olympic Road Race problem is modeled on the Cycling Road Race event in the Summer Olympics. The most significant aspect of this very long race is that a cyclist riding behind another expends up to 30% less energy by "drafting" behind another rider, as the rider in front absorbs much of the wind resistance. In this problem, we were to design a "coach" for a team of riders. Each team could consist of as few as one rider, or as many as ten. A rider starts the race with a specific amount of energy, and energy consumption throughout the race is $v^{2.5}$ units/sec, where v is the speed of the rider. However, in the case where the rider is some distance d behind another, the energy consumed per second is reduced by a factor $f(d) = (5-d)/10$ for $2 \leq d \leq 5$ and 0 otherwise. Riders can accelerate or decelerate at a rate of up to 1 m/sec. However, no rider may exceed 25 m/sec. A team scores points as follows: 5 points for a gold medal, 3 points for silver, and 1 for bronze. The objective, of course, is to place in the race and score the most points.

In this paper, we will first analyze the challenges presented by this problem and some of the strategic issues discussed in class in the section "Problem Analysis". We will then detail the strategy we chose to develop and its implementation in the "Player Implementation" section. Here we will also discuss some of the problems we encountered in implementing our player. Finally, we will analyze the class tournaments in "Tournament Results".

Problem Analysis

The Olympic Road Race problem features some interesting issues. In developing an effective player, one must consider the cost of movement relative to speed, the variety of general team strategies and the constraints on implementing them, and offensive and defensive maneuvers against those strategies.

Cost of Movement

The important aspect of the cost of each move in Olympic Road Race is the fact that it is not linear relative to speed, but is $v^{2.5}$. That is to say, the faster a rider travels, the more energy per unit speed it consumes. Because of this, it is disadvantageous to vary a rider's speed widely during the race. The average speed is not as important as maintaining a fairly constant speed throughout the race.

General Team Strategies

"Line" Strategy

The most popular strategy discussed in class, and the strategy we chose, as we found it the most effective, involved all the players of a team aligning themselves behind one another at the beginning of the race and traveling together at a constant speed. This would give the player in the back (the “sprinter”) the most benefit from drafting behind the others. Call this the “Line” strategy. The idea is that the leader of the team would die at some point, leaving a new leader who had benefited from drafting behind the first. This process would be repeated until the sprinter was left to ride alone for the last leg of the race, using its last bit of energy to cross the finish line. Because the percentage of energy saved in drafting is 30%, each leader would die after leading for approximately 70% of the distance from the time it became the leader to the finish line, varying slightly with the number of team members.

This strategy divided the problem into two main stages: aligning the riders and traveling and maintaining the group. One of the challenges in the second stage was to find a formula for the appropriate speed the team should travel. Several equations were discussed, and they will be addressed in the section “Player Implementation”. These formulas, however, produced in the first stage another challenge: to minimize the time and energy spent to align the team members. Since many of the teams using this strategy were using similar velocity formulas, it was often the efficiency of this first stage that made the difference between a team who won the race and one who did not. Our player tried three first stage strategies, which are also discussed in “Player Implementation”.

Other Strategies

While many groups chose to implement the Line Strategy, some took an “Offensive” or “Leech” approach. The Offensive strategy was intended to thwart the performance of the better players, who often used the Line strategy. The Leech strategy, on the other hand, was intended to benefit from the good performance of these better players.

In the Offensive approach, a team would sacrifice *R-1* riders as “blockers”. The blockers would speed ahead of the other teams, position themselves in front, then slow down, forcing the other teams to slow down as well. The remaining player would simply travel at a constant speed toward the finish line. There are two main drawbacks to this approach. First, the number of teams *T* is often be greater than the number of players *R*. Because of this, the *R-1* blockers could not possibly block all the other teams, leaving the possibility that the best teams perform unhindered. Second, the single rider who was intended to finish the race would never get the benefit of drafting.

Teams who implemented the Leech approach usually sent each rider on its own to position itself behind another team. The rider would draft behind this team throughout most of the race, conserving its energy. It would then pull around and pass its host team at the last moment, hoping to win the race. The primary difficulty in this strategy was to choose the right team to follow, as drafting behind a team with a poor strategy would produce a situation only slightly better than poor.

Both the Offensive and Leech strategies posed some problems for every team. To counter these approaches, many teams implemented a defensive “wiggling” maneuver, so as to try to shake the player leeching or blocking. The initial wiggling implementation discussed in class involved a host team moving to the right or left with 1/2 probability. This gave the leech rider the opportunity to maintain its position behind the host team with a probability of 1/2. It was decided that a better implementation might be to move to the right, left, or stay in the same lane, each with probability 1/3, leaving the leech only 1/3 probability of maintaining its position. While wiggling didn’t entirely prevent other riders from drafting, it would generally allow the host team to lose the leech for enough of the time to prevent the leech from gaining a winning benefit.

Player Implementation

Stage 1: Getting Together

For those teams using the Line strategy, the difference between winning and losing is the speed at which the line of riders is formed. Since we implemented this strategy over the course of the past three weeks, we have been constantly trying to tune our line-up method. We found this step to be difficult, as the acceleration calculations complicated the process of knowing when to speed up and slow down in time to move into line. The overall idea of our implementation is to lose as little of the sprinter's energy as possible and line up the teammates in accordance with their energy levels. So the rider with the least energy would be at the front and the rider with the most energy would be at the back.

We attempted many different ideas, but only implemented three for in class trials and the tournament.

Initial strategy

Our first idea was specifically developed for the idea of having the rider at the back have the most energy. The first step is to choose a meeting lane. We take the average of the initial lane numbers where the riders are placed, and this is the designated the meeting lane. All of the riders start moving toward this lane. The rider closest to the meeting lane becomes the rider at the back of the line. The next closest to the meeting lane becomes the next from the back and so on. See Fig. 1 and 2 for an example.

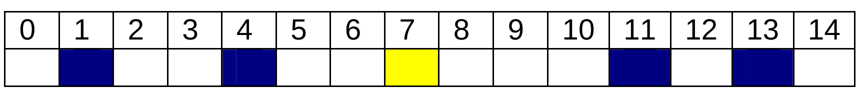

Fig. 1: In this configuration the meeting lane is lane 7. Thus the rider in lane 4 is the sprinter, 11 is next, 1 is after 11, and 13 is the leader.

Start Line
4
11
1
13

Fig.2: Here the riders from Fig.1 are shown after they get into formation.

The first step is to move the riders towards the meeting lane and then get them into their positions. Because it gave us the most control over the riders' positions, we accelerated each rider by 1 if it wasn't already in its position relative to the intended sprinter, and then gave each an acceleration of 0 until it got to its position. Once a rider was in its proper position, we gave it acceleration –1 to slow it to a stop. The players waited for all other teammates to move into position before moving to the next stage of the race. This scheme lined the team up perfectly, 2 meters apart, with *2 * position* energy loss.

This strategy was very good for moving them into formation expending the smallest amount of energy, but it was not very effective for winning the race against other players for two reasons. First, we found that other riders easily got in the way, delaying progress, if not blocking our team entirely.

Improvement

An improvement on our original strategy was to have all riders race to 10 meters down the track and then move into the same formation. We found that this change made collisions at the starting line less of a problem for our riders, and we were winning some offline tournaments. However, stopping in order to line up our players was still leaving us at a disadvantage against the black and red teams who were using a similar race strategy. We found that waiting was making us lose too much time in the race, and we could not regain it with a higher velocity. Therefore, we decided to explore other options.

Final Strategy

We thought about how we could merge our riders into a line while moving at a constant velocity. We tried accelerating the leader to the desired velocity, then moving the other riders into their assigned positions. We implemented this and found that one rider would often get behind the leader early, leaving any rider whose position was to be between them no chance of getting into place, because they would accelerate to catch up, then decelerate too late to fall into place. We decided that this problem might be solved if the leader would not speed to the full desired velocity, but half of the speed. With each rider going half the desired velocity, we did not have the problem a rider exceeding the desired speed when trying to move into position.

The riders begin by accelerating to half the desired velocity. Once they are up to speed, they are all at the same meter marker, and they start to figure out who should slow down. Each rider knows which position they are supposed to be in. The rider behind the leader (position *R-2*) must decelerate by 1 for one second, then move at its new velocity for one second, at which time it would be 2 meters behind the leader. Then it must accelerate by 1 again to move back up to the speed of the leader and move into position. The rider in the next position would decelerate by 1 for one turn, move at this new velocity for 3 turns, then accelerate by 1 to move into position 4 meters behind the leader. Each player other than the leader, therefore, decelerates by 1 for one second, moves at the new speed for *2(R–position number–1) – 1* seconds, then accelerates by 1 once again to move over and fall into position. Fig. 3-11 show the progression of this tactic.

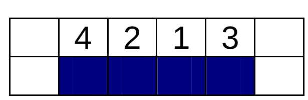

Start Line

Fig. 3: Players have reached ½ desired velocity at time step (TS) *x*. Numbers show assigned positions, each block is 1 meter.

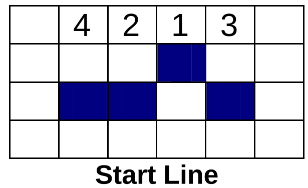

Fig. 4: TS = x + 1: Riders 4,2, and 3 with acceleration –1 m/s

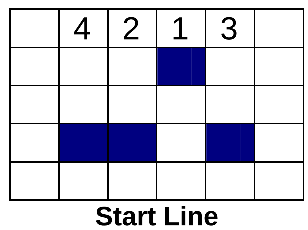

Fig. 5: TS = x + 2: Riders 4,2, and 3 with acceleration 0 m/s

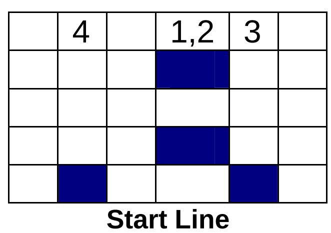

Fig. 6: TS = x + 3: Riders 4 and 3 with acceleration 0 m/s; Rider 2 with acceleration 1 m/s

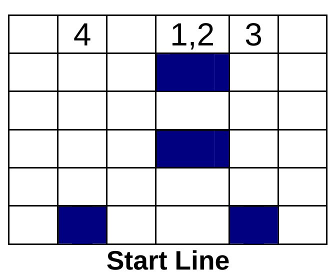

Fig. 7: TS = x + 4: Riders 4 and 3 with acceleration 0 m/s; Rider 2 with acceleration 0 m/s, in position

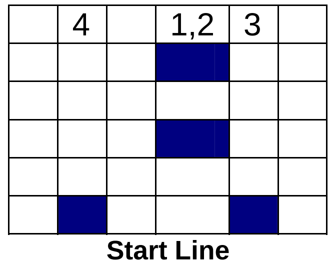

Fig. 7: TS = x + 5: Rider 4 with acceleration 0 m/s; Rider 3 with acceleration 1 m/s

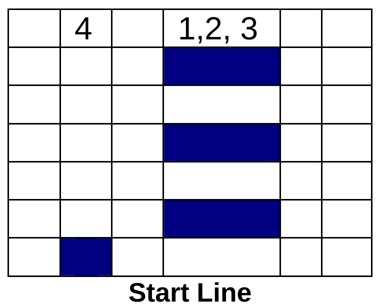

Fig. 8: TS = x + 6: Rider 4 with acceleration 0 m/s; Rider 3 with acceleration 0 m/s, in position

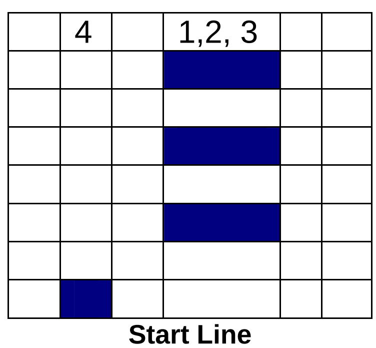

Fig. 9: TS = x + 7: Rider 4 with acceleration 0 m/s

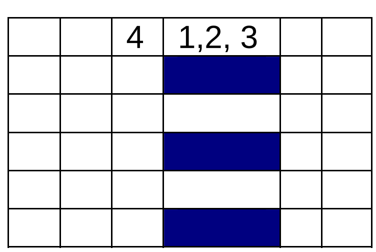

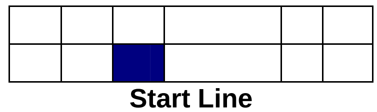

Fig. 10: TS = x + 8: Rider 4 with acceleration 1 m/s

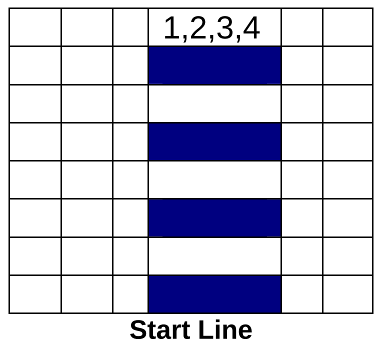

Fig. 11: TS = x + 9: Rider 4 with acceleration 0 m/s in position

This tactic requires that all riders accelerate to a velocity equal to *floor(desired velocity/2)*, then run the above described algorithm to get into line. Therefore, this tactic takes *floor(desired velocity/2) + 2R – 1* turns. We found that this was efficient for getting our riders into a line without losing too many turns in waiting for each rider to move into position and while still making progress in the race.

Stage 2: Velocity Equations

Several different equations were introduced at the start of this project. Though our equation is identical to Vipul's, we derived it from scratch when it seemed that there were many different equations in use, all yielding a slightly different speed. To arrive at our velocity equation, we simplified the problem down to two players and then expanded it to n.

The starting point for all of the equations is the simple "max speed equation." This should yield the optimal velocity for a single rider with a given amount of energy to just finish a race of length d.

$$v = \left(\frac{E}{D} \right)^{\frac{2}{3}} \tag{1}$$

Expanding this to two riders is relatively simple. We need to divide the race length into two sections, with the first rider dying at the end of the first section. To do this, we set *x* to be the proportion of the total distance the first rider covers before dying. E_0 and E_1 are the starting energies for rider 0 and 1, respectively.

(2) $$\frac{xD}{\left(\frac{E_0}{xD}\right)^{\frac{2}{3}}}+\frac{(1-x)D}{\left(\frac{E_1-.7E_0}{(1-x)D}\right)^{\frac{2}{3}}}$$

The second rider's velocity calculation relies on the first rider's energy. If we were to expand the field to three riders, the last rider would depend on energy from the last two. This creates a recursive definition:

$E_{0f} = E_{0i}$
$E_{1f} = E_{1i}-.7*E_{0f}$
$E_{2f} = E_{2i}-(.7E_{1f}+.7E_{0f})$
…

(3) $$E_{nf} = E_{ni} - \left(\sum_{i=0}^{n-1} .7E_{if}\right)$$

This can be visualized in the chart in Fig. 12.

Energy Expended	Segment 0	Segment 1	Segment 2
Rider 0	E_{0f}		
Rider 1	$.7E_{0f}$	E_{1f}	
Rider 2	$.7E_{0f}$	$.7E_{1f}$	E_{2f}

Fig. 12: Stages of the energy throughout the race.

Now, since in each segment, the rider below uses only 70% of the energy of the rider above (if the rider is trailing someone), we can look at it as if the rider below is gaining 30% of the energy expended above him to his total energy. It follows that in the case of two riders, the second rider's total energy is equal to his initial energy plus 30% of the first rider's energy, assuming that $.7E_{0i} < E_{1i}$. If we expand this to one more rider, we can consider the first two riders one complete system:

$E_{0i} = 100$
$E_{1i} = 100$
$E_{2i} = 100$
$E_{1total} = 130$
$E_{2total} = E_{2i} + E_{1total}$

Therefore:

(4) $$E_{2total} = E_{2i} + .3E_{1i} + .3^2E_{0i}$$

Our final energy summation may be generalized to:

(5) $$E_{nTOTAl} = \sum_{i=0}^{n} (.3)^{n-i} E_i$$

And the final speed equation for the race after plugging into (1) is:

$$v = \left(\frac{\sum_{i=0}^{n} (.3)^{n-i} E_i}{D} \right)^{\frac{2}{3}} \quad (6)$$

This velocity is calculated once as soon as the riders are in a line. This ensures that they are drafting correctly and that the benefits are exactly 30%. The speed does not change throughout the race (if the equation is correct and nothing is interfering with the players), and therefore it only needs to be calculated once unless a speed-affecting event occurs, such as a blockage.

The equation is fudged slightly in the beginning to ensure that the player does not die right before the finish line. A tiny fraction of the course length is added to the total distance to reduce the speed very slightly. This compensates for general rounding errors and game imprecision.

To further compensate for possible errors, the velocity is recalculated using equation (1) upon the death of the second-to-last rider. This ensures that the rider will generally at least finish the race if a blockage occurs. We decided against calculating the velocity at every turn because it gave us no perceivable advantage—in most races 3-4 turns is the difference between winning and losing a race, and if a serious blockage occurs, the player is almost definitely out of the running.

We found this equation to be very reliable. This is evidenced by the final speeds reported by the last riders after the recalculation matching the speeds generated at the beginning of the race.

Defense against Enemies

To compete in a multi-player environment, our player utilizes two separate, but simple techniques. Once it is in a line, it attempts to switch lanes at every turn. To ensure that the line remains intact, it first verifies that the space (with an added buffer to prevent blockage) to its side is clear of enemies. This prevents a Leech rider from realizing a benefit by trailing us.

The second attempt at multiplayer competition is a rudimentary trailing strategy. After 200 turns, if our player is not the leader, and the leader is in a line and not swerving (verified by keeping a moving average of the leader’s lane changes), our player attempts to trail the leader to win the race at the margin between the final rider and the finish line. Since most players utilize the swerving strategy, we found that this strategy is rarely triggered.

Tournament Results

Our starting get-together strategy resulted in very consistent results, but unfortunately also slowed us down a small number turns, which often meant that we were not one of the top three finishers. Our consistency is exhibited in the “two of the same team” tournament results. Predictably, we finished with an average score of roughly 4.5. However, the players that led the default tournaments usually finished below us. This is likely because their players would conflict with one another.

We also placed in the top third (in average score) for the default conditions with 9 riders and had the second most number of wins. Additionally, we placed 4th in the race with a short length. These two results show that our player is able to survive rough conditions and has a good starting out strategy.

We think that our player would perform well in particularly rough conditions, such as a large number of riders. Our strategy of aligning players in order of energies would prove very successful in a situation where precision is significant.

Group Participation

Alex Davidov developed correct velocity calculations and rider parasitic behavior.
Becky Plummer developed initial velocity calculations and getting together strategies.
Anya Robertson worked on player precision improvement and developed final getting together strategy.

Acknowledgements

We would like to acknowledge Bogdan Caprita for volunteering the first velocity equation calculating the optimal velocity to get the sprinter across the finish line. We used Bogdan's equation until Vipul Bhasin volunteered a second equation that the class agreed was the best equation for the solution. Using these velocity equations, we implemented the Line strategy as discussed by the class.