

Optimal gear selection on an automatic bicycle

by *Iain Crouch*

Abstract

This paper describes an electronically controlled automatic transmission which was designed and built for use on a bicycle as a final year university project. The particular focus of the project was on the algorithm which determines the optimal gear to select at any time for maximised performance; i.e. acceleration and top speed, given the amount of effort the cyclist is putting in. This optimal gear selection is shown to be possible with the correct use of a prior strategy based on a fixed target cadence (pedalling rate). However, this cadence is unknown and varies over time and between different cyclists. It is shown that a cyclist's optimal cadence can be continuously estimated during normal cycling by a controller which fits recorded data to an assumed model. Microprocessor-based hardware was constructed and fitted to a bicycle to allow such a controller to be implemented and tested, and examples of the results are given and discussed.

Motivation

Automatic gearboxes have been fitted to motor vehicles for decades, yet rarely feature on human powered vehicles. The technical problems are obvious: the extra weight, cost and complexity of an automatic governor would be significant when fitted to a simple, lightweight cycle. In addition, existing 'crash' transmissions are not designed to transmit power during gear changes; to modify them would further compromise weight, strength, cost and efficiency. Using an electronically controlled gear selector can help reduce the weight penalty, and such devices have been introduced as optional extras on continuous-power transmissions such as the Shimano Nexus Auto-D hub gears and the Browning split chainring system (Kyle, 1995). Mechanical compromises are still present, however, so they are aimed in general at recreational cyclists - the Shimano system is specifically targeted at novice cyclists who find manual shifting complicated or distracting.



Existing automatic systems usually employ a simple gear selection strategy that attempts to maintain the cyclist's cadence at some fixed value, say 64 rpm. This is a comfortable rate for most people, however it is limiting in terms of performance – for example a racing cyclist would develop his peak power at a much higher rate. Furthermore, the selection of the best gear for maximum acceleration depends on characteristics that vary between cyclists and over time (due to fatigue for example). Some systems (including the Browning automatic) have the facility for training or adapting to the cyclist's preferences, which will improve performance but is still not optimal as it relies on the perception of the human rider.

The focus of this investigation was on improving the 'intelligence' of the automatic gear selection strategy, to assess whether more experienced cyclists could then find performance gains in using automatic rather than manual selection. To accomplish this, the extra intelligence must not only make up for the additional weight and complexity, but also the disadvantages of taking control away from the rider – for example a human cyclist has the advantage of being able to see changing conditions ahead. However (s)he also takes the effort of performing the gear change into account when considering changing, and will often only change gear



Figure 1, top: Modified transmission, showing automated derailleur, torque sensor and reed switches

Figure 2, below: Main unit

when (s)he is uncomfortable with the gear (s)he is in, whereas an automatic selector may always choose the optimal gear given the information that it has.

It will be shown that a cadence-based selection strategy can be optimal if it uses the correct target cadence.

Hardware Description

As a university project rather than a commercial development, there was the opportunity to disregard some mechanical difficulties and practical issues and concentrate on investigating the control aspects. The prototype used to develop and test the control algorithms was therefore based on a conventional 'crash' gearchange, with modifications built for durability rather than light weight. This means that, although the drivetrain components were new and therefore as smooth-shifting as they could be, any

cyclist using it must be vaguely aware that unexpected gearchanges would occur. The practicalities of commercial implementations on other drivetrains are considered in the 'Discussion' section.

The test bicycle was built around a Giant Terrago MTB frame, fitted with a Shimano Deore LX groupset and slick tyres. Figure 1 shows the modified transmission, which has a 9-speed cassette and a single front chainring. The existing derailleur was modified using a geared motor and a feedback potentiometer, which allows its position to be sampled by the onboard 8-bit PIC microcontroller and compared with the position for the desired gear. The motor was a surplus component whose specification far exceeded the estimated requirements; it rotated at 8rpm and produced 0.6 Nm (1.8 Nm peak) torque - greater than the torque from the derailleur's original return spring. Despite the motor's high inertia, a simple proportional feedback control routine, combined with the natural damping of the system, gave a satisfactory response for the derailleur, with a very slight overshoot to aid shifting.

Two magnets were mounted on opposing rear wheel spokes, allowing the microcontroller to time pulses from reed switches mounted on the bicycle's chainstays and calculate its speed. Similarly, magnets mounted on the cranks with corresponding reed switches on the chainstays send pulses to the microcontroller when the cranks are approximately horizontal. A chainring (actually the 'granny ring' from the bicycle's chainset) was mounted on a machined disk, which in turn was attached to a cantilever arm via a bearing and held against the tense upper part of the chain. The applied torque could then be measured using a pair of strain gauges fixed to the cantilever. The signal is amplified and sampled by the microcontroller's analog to digital converter. Data is regularly logged so it can be downloaded to a PC after each test run. The control system is completely self-contained, with batteries and a display built into a main unit (figure 2), and control routines, written in assembly language, running on the microcontroller.

The Model Used

To allow investigation, simulation and development of the control strategy, a general model of cyclists' output characteristics was required. A parabolic approximation to the relationship between maximum power and cadence was found to be satisfactory: derived power versus cadence curves [including Whitt and Wilson, 1974] were found to closely approximate parabolas, especially over the region of interest close to the cyclist's optimum peak power cadence. The relationship is shown as a graph in figure 3. As this parabola is known to pass through the origin, the model of available power P versus cadence ω of a cyclist may be completely defined by two values: his or her peak power P_{MAX} and corresponding optimum cadence ω_{OPT} :

$$P = P_{MAX} (1 - (\omega - \omega_{OPT})^2 / \omega_{OPT}^2)$$




Figure 3: Diagram of the relationship between a cyclist's maximum power output and his pedalling rate

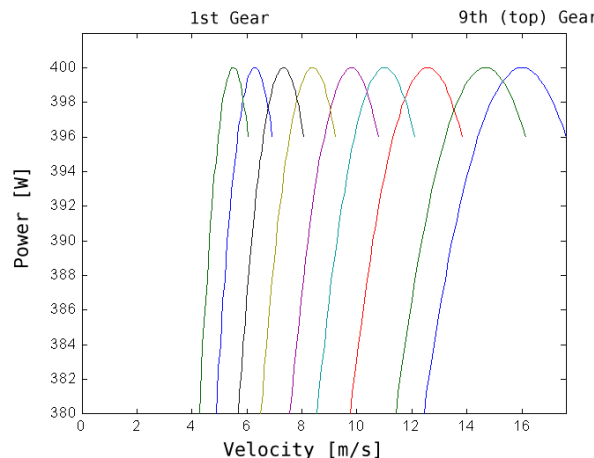


Figure 4: A Matlab graph of the corresponding available power versus road speed for different gears.

The cyclist's available power versus road speed, mapped through the available gears, is shown in figure 4. The graph




Figure 5: Torque vs. Pedalling Rate

shows that correct gear selection is essential to maximise performance – an automatic controller can calculate this correct gear, given the road speed, gear ratios and the cyclist's optimum cadence (his peak power need not be known, as this only scales the graph vertically). While road speed can be measured accurately, and the gear ratios are known, the optimum cadence remains unknown and variable. Running simulations (in Matlab) of maximum rate accelerations indicated that errors in the target cadence value used by a controller of just 10rpm would noticeably restrict performance. The variation in efficiency between gears is negligible and was disregarded.

Control – estimating the optimum cadence

This parabolic model not only allows a mathematical analysis of the effects of automatic gear selection, but also provides the key to allowing continuous online estimation of the optimal Cadence ω_{OPT} the one missing parameter that is required for theoretically optimal control.

A parabolic power versus cadence relationship implies a linear crank torque versus cadence relationship. This was verified using collected data such as the maximum rate acceleration in figure 3. The zero-torque intercept of the fitted line corresponds to the zero power intercept of figure 3; hence the cyclist's peak power is found at half the maximum cadence predicted by the fitted torque data. (NB this will not necessarily be the cyclist's actual maximum as it is outside the range of the data used to develop the model). This is the basis of the optimum cadence estimation algorithm - in figure 5 the intercept of 164rpm predicts, plausibly, an optimum cadence of 82rpm. Furthermore, data for other, less than maximum rate accelerations was also found to approximate straight lines, albeit less

cleanly, with correspondingly lower estimated optimal cadences. This was an early indication that the controller would be capable of adapting to cyclists when they are not fully exerting themselves, despite the less well-defined nature of the corresponding data.

Test data sets of torque/cadence pairs over typical accelerations were taken and used to develop the estimation algorithm, which was then implemented and tested on the prototype bicycle in the following form:

A single torque sample is taken every right crank stroke (for simplicity, and consistency of data) and recorded along with the current cadence. This then updates an array which contains the last eight torque/cadence data pairs. This array is sorted to detect outlying points (bad data) and allow the gradient of the torque/cadence relationship to be approximated using a specially designed regression technique. The new optimal cadence estimation can then be found by extrapolating the line to find the cadence at zero torque, and dividing it by two. The previous target cadence is then updated by averaging it with the new estimation, which is weighted (or rejected) according to the conditioning of data (a better spread will give a better approximation) and how well the data points fit the linear approximation (best if the cyclist is fully and continually exerting himself). Other criteria were imposed to simplify computation, for example the cyclist must be accelerating (as data collected when cycling at constant speed is ill conditioned, while usable data from decelerations is rare).

The new target cadence is then used by a gear selection routine to calculate the range of cadences appropriate to the current gear. It can then shift up or down accordingly if the current cadence is outside the range. Only one gear change is permitted per crank revolution to avoid damaging the actuator.

Results

Although exhaustive testing is beyond the scope of the investigation and unnecessary at this stage of development, data collected over several tests provides convincing evidence that:

- the estimated optimum cadence coincides with the target cadence for best performance
- it is capable of tracking changes in optimal cadence

- the response is stable, robust, consistent and fast, surpassing the early aims of the project




Figure 6: Timed sprints for fixed cadence controllers

Figure 6 was generated by timing sprints, with the controller using a fixed target cadence – 2 times were recorded at each multiple of 10rpm. The graph indicates that the cyclist's performance is maximised if his cadence is kept close to 100rpm. An early version of the estimation algorithm was also running; although the runs were short to avoid tiring the cyclist (<200m), it still regularly estimated optimum cadences close to the actual optimum of 100.




Figure 7. Evolution of estimated optimum cadence over a test run

Figures 7 and 8 are examples of logged data, independently taken for test runs using the estimation algorithm. The data was recorded over urban courses, providing rich information for the controller due to frequent accelerations after corners and junctions so that its

behaviour could be observed. In figure 7, the cyclist is fully exerting himself for the duration of the run, and is therefore exhibiting a well defined and relatively constant power-cadence relationship. The target cadence starts at 60 by default, but converges smoothly and rapidly to 74 during the first hard acceleration. The single minor update to 75 during the next acceleration is further evidence of convergence. In addition it suggests not only that the controller is capable of extracting information from further accelerations but that the estimated optimum found is close to the previous estimate, which was calculated from an independent data set. This implies a degree of consistency in both controller and cyclist behaviour. The reasonably high estimated optimum of around 75rpm corresponds to the cyclist's effort level.

Figure 8 was recorded by a cyclist who was not fully exerting himself, yet the target cadence still converges rapidly over each acceleration, from the initial value of 80 (the flat data between 33 and 66s was caused by the chain coming off!). The final, maximal exertion sprint provides the clean torque data to allow the target cadence to jump rapidly from 64 to 94rpm. Although the investigation was only intended to demonstrate the ability of a controller to converge on an optimum slowly, assuming continuous exertion and slow rate of change of ω_{OPT} due to fatigue, even this simple incarnation backs up the earlier suggestion that the controller would be capable of tracking ω_{OPT} at sub-maximal effort.

Other sets of test data over different conditions show similar characteristics; no irregularities were ever observed in the target cadence found, but the fast response means that there is scope for further filtering to reduce sensitivity if desired.

In the opinion of the users, the gears selected by the controller felt natural and comfortable, especially as the controller shifted more often than they would normally bother with. By contrast, if the target cadence of the controller was fixed,




Figure 8. Evolution of estimated optimum cadence over another test run

users found the enforced cadences to be restrictive, forcing them to cycle more leisurely or aggressively than they desired.

Discussion

An analytic approach to optimising automatic gear control on HPVs has been investigated, developed and demonstrated. The resulting controller is capable of adapting to the changing rider characteristics despite requiring no prior training or setting up; the user does not have to intervene or even be aware of the controller's operation. Although a thorough evaluation is not possible at this stage due to the nature of the prototype transmission and the amount of variation in characteristics that would have to be accounted for, the investigation provides convincing support for the theory on which it is based. This suggests that the algorithm is indeed worthy of further development for more suitable transmissions.

The parabolic power / linear torque model used is surprisingly simple and effective, and lends itself well to the two stage optimum cadence estimation and gear selection algorithm. Other approaches based on the same theory are possible – for example a controller could aim to minimise steps in the torque applied at the wheel over gear changes. However the approach used has the advantage of also using the model to filter the input data, at the linear regression stage. Success with the analytic approach meant it was not necessary to resort to common

learning techniques (such as neural networks or fuzzy logic), however for example fuzzy rules could be used to add heuristics, such as a cost associated with changing gear.

Gear selection based on prior knowledge of theory and accurate measurements, rather than the cyclist's own preferences, habits and perception, should result in performance gains. In the case of this project that means faster acceleration, and although a direct comparison is not possible

due to the crash transmission, the results obtained are a strong indication that the advantages of automatic gear selection can outweigh the disadvantages for cyclists at all levels. The system has great potential for further development to achieve this: the control software uses only a fraction of the simple 8-bit processor's time, so there is much scope for increasing the complexity and flexibility of the controller. The main unit and torque sensor can be made very light, especially if the torque sensor is moved to the crank itself and the actuator is made more efficient to reduce battery requirements (the average current at the moment is a few milliamps). Cadence and crank position can also be determined from the torque variation. Weight is also saved in other areas due to the removal of the mechanical derailleur cable, shifter and return spring. Having more precise control of the derailleur position and shift timing may allow the crash transmission to be adapted for faster, continuous-torque gear changes. Furthermore, the more frequent and correct gear selection may mean that fewer gears are required, possibly with a single front chainring to avoid extra automation in the case of the crash transmission.

The optimum cadence estimator could find other uses which do not depend on developing an appropriate transmission. For example it could be the basis of a training aid, which would resemble a

cycle computer with a remote crank-mounted torque sensor, that could tell a cyclist which gear to be in at any point as well as recording his optimum cadence variation over time.

Conclusion

This investigation has resulted in the development of simple basis for electronic gear selection which merits further investigation and has the potential to benefit riders of even the highest standards. At the very least, it has enlightened the author on the nature of cycling in general and improved his own gear selection greatly.

References

- Whitt, Frank R, Wilson, David G. 1974. *Bicycling Science*. MIT Press, Boston
- Kyle, Chester R. 1995. *The Browning Automatic Bicycle Transmission*. *Cycling Science Winter 1995*

The author

Iain Crouch is a mountain biker from Glasgow, Scotland. He has recently completed an MEng in Engineering Science at Oxford University and can be reached at: <iaincrouch@hotmail.com>

Letter: recumbent monocycle

I have been searching for pics and/or drawings of a pedal powered monocycle. (you ride inside the wheel...usually four to five feet in diameter.) I would like to know if you have any info/links or other news. My next project will be a pedal boat using a hydrobike drive unit. I recently raced a full size one for three miles and came in halfway in the pack amongst kayakers and canoes. This is a resounding success given my previous last place finishes with a home built chain-drive unit! A picture of what I want to build can be seen at: <<http://www.valedo.com/5370410.gif>>

Brian Burgess
<pedalyurassoff@peoplepc.com>

[Any reader with more info please contact Brian Burgess.. His link is pictured below. Ed.]

