OFFICIAL JOURNAL OF THE BRITISH HOROLOGICAL INSTITUTE

The Horological Journal

NOVEMBER 2022
www.bhi.co.uk

A New Free-Pendulum Clock

Synchronising Without a Synchroniser

Simon Goodman

Free-pendulum clocks are fascinating machines, and I was inspired particularly by Philip Woodward's clock, W5, to build my own. Like W5, it had to be entirely mechanical.

A free pendulum does not interact with its escapement except to receive impulses to maintain its oscillation. Many attempts have been made to design free-pendulum clocks; those of William Hamilton-Shortt, **Figure 1**, were the most successful. Philip Woodward re-interpreted Shortt's system for the domestic environment, and his 'W5', **Figure 2**, managed to maintain the theoretical purity of the original.

The Shortt clock has a free (master) pendulum and a slave pendulum, the latter being a slightly modified Synchronome. Both are impulsed by gravity arms. The impulses are applied by releasing a gravity arm (pivoted in the clock movement), which then falls onto an impulse arm (attached to the pendulum rod). In the case of the master, there is a delicate roller mounted on the tip of the impulse arm. In the slave clock, the roller is fitted to the gravity arm, acting on a curved face of the impulse arm. In each case, the pendulum receives its impulse by the gravity arm falling off the impulse arm, pushing the pendulum as it swings away. When the gravity arm falls free, it is electro-magnetically raised and reset for the next impulse.

Figure 3 outlines the impulse sequence for the Shortt clock. A jewelled gathering arm on the slave pendulum rotates a count wheel, taking 30 seconds. This releases the slave's gravity arm to impulse the slave. When the gravity arm falls free of the slave's impulse arm it closes an electrical contact, energising an electro-magnet to reset the slave's gravity arm and to send an electrical signal to release the master pendulum's gravity arm. The master's gravity arm impulses the master pendulum; when the gravity arm falls free of the master pendulum's impulse arm, it closes an electrical contact that energises an electro-magnet to reset the gravity arm and to send an electrical signal to trigger a 'hit-and-miss synchroniser' to bring the slave into phase with the master.

The master pendulum controls the rate of the overall system, while the slave pendulum times the impulses to the master. The slave's period will inevitably drift from the master's over time, and so the slave must be corrected occasionally as the clock runs to keep it in phase with the master. This is done by the hit-and-miss synchroniser. In practice, the slave is set to a slightly losing rate; the synchroniser periodically advances it to keep it in phase with the master.

The synchroniser consists of two parts. A synchroniser spring in the form of a stiff spring-steel strip is mounted vertically to the slave pendulum rod. Its lower end is attached to a block on the pendulum rod, but its upper end is free. Perpendicular to this is a metal blade capable of being pulled downward by a nearby electro-magnet. When the magnet is energised, the blade is momentarily pulled into the path of the swinging pendulum synchroniser spring. If the slave's phase is

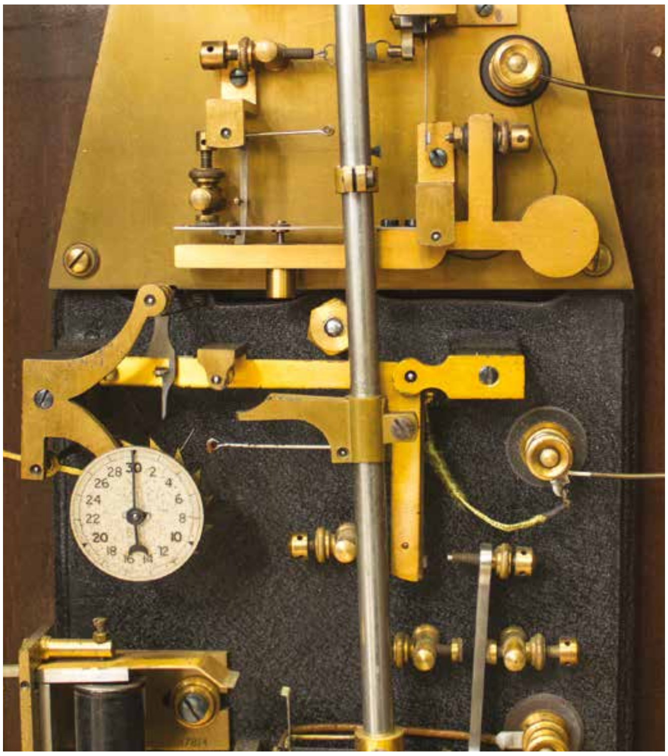

Figure 1. Detail of the Shortt slave clock.

Figure 2. Philip Woodward's W5.

behind the phase required to impulse the master on time, the blade ‘hits’ the spring, applying a force towards the centre of the slave’s swing, advancing the phase of the slave (speeding it up). If the slave’s phase is ahead of the phase required to impulse the master on time, the spring will already have passed the blade when the synchroniser is triggered and so ‘misses’ the spring.

W5 follows a similar series of steps to the Shortt clock, **Figure 4**, but all of them are entirely mechanical and the operation is elegantly simplified by using the same gravity arm to impulse both pendulums, acting first on the slave, and then on the master. After the gravity arm falls from the master’s impulse roller, it is mechanically reset and a hit-and-miss synchroniser is triggered to correct the phase of the slave.

So, in both the Shortt and W5 clocks, the end of the slave’s impulse starts the master’s impulse, and at the end of the master’s impulse, feedback is sent from the master to the synchroniser to adjust the phase of the slave.

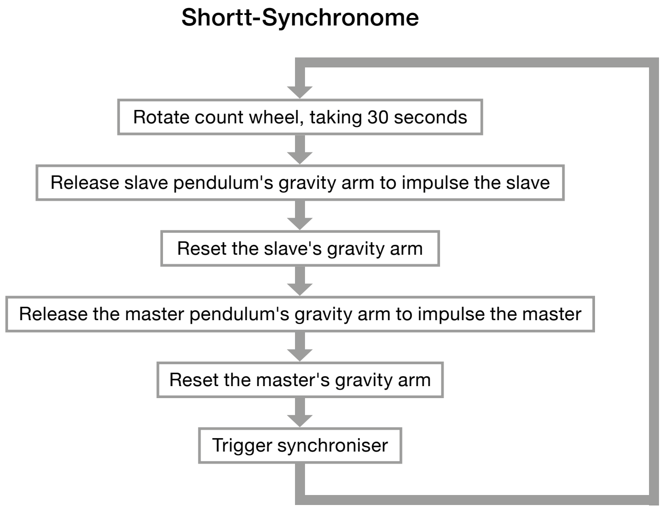

Figure 3. Shortt clock impulse sequence.

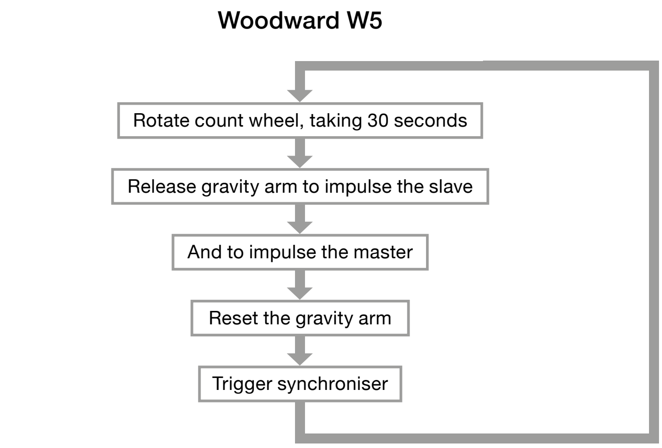

Figure 4. W5 impulse sequence.

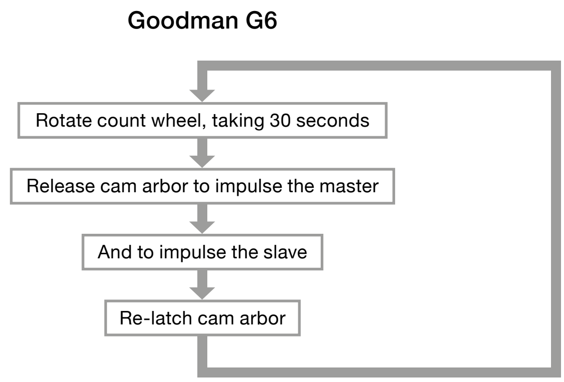

Figure 5. G6 impulse sequence.

In Philip Woodward’s book *My Own Right Time*, he likens an escapement driving a free (master) pendulum to a servant bringing the master his tea at the time when the master wants it, without the master having to ask for it. However, in both the Shortt and W5 clocks, it seems to me that when tea time arrives, the servant first drinks his own tea (by impulsing the slave) before delivering tea to the master.

In my clock, the servant (much more appropriately) delivers the master’s tea *before* drinking his own. This simplifies the escapement enormously, including completely eliminating the synchroniser.

My clock uses free (master) and slave pendulums, and **Figure 5** shows its impulse sequence. The slave drives a count wheel to trigger an impulse to the master every 30 seconds. The end of the master’s impulse triggers the slave’s impulse. The slave is synchronised in every impulse sequence, but there is no synchroniser. Instead, the slave impulse mechanism automatically introduces a variable (synchronising) escapement error into the slave, which depends on the phase of the slave when its impulse starts.

An escapement error can advance or retard the phase of a pendulum, and if the force applied to the pendulum by the escapement acts towards the centre of the pendulum’s swing, the escapement error advances the phase of the pendulum. This is the arrangement I have used.

Figure 6 (**overleaf**) shows the key components of the escapement. I should perhaps name this clock G6 as it is the sixth clock I have built.

The master pendulum **1** is suspended directly in front of the shorter slave pendulum **2** (the master pendulum’s suspension and bob are not shown because they are above and below the part of the clock shown in **Figure 6**). Each pendulum carries an impulse roller **3**, **4**. The escapement comprises a cam arbor **5** carrying cams **6**, **7** for applying impulses to the impulse rollers. A constant clockwise torque is applied to this arbor.

The slave pendulum carries a gathering pusher **8** for advancing a count wheel **9**, positioned above the cam arbor. A back-stop prevents the count wheel rotating clockwise.

The cam arbor carries a latching wheel **10**, which has four latching surfaces at 90° intervals, for engaging with a pallet **11** on a release lever **12** to stop the cam arbor rotating clockwise. The release lever pivots at its lower end, below and to the left of the cam arbor, and extends upwardly to the left of the latching wheel. One of the count wheel teeth is cut deeper than the others so that once in every rotation of the count wheel the pusher falls into the deeper tooth, contacts the upper end of the release lever and moves it to the left, unlatching the cam arbor and freeing it to rotate 90° until it latches again.

The master’s drive cam **6** and the slave’s drive cam **7** each have four lobes at 90° intervals, so that an impulse sequence takes place each time the cam arbor is unlatched, and the cam arbor makes a complete rotation once every two minutes.

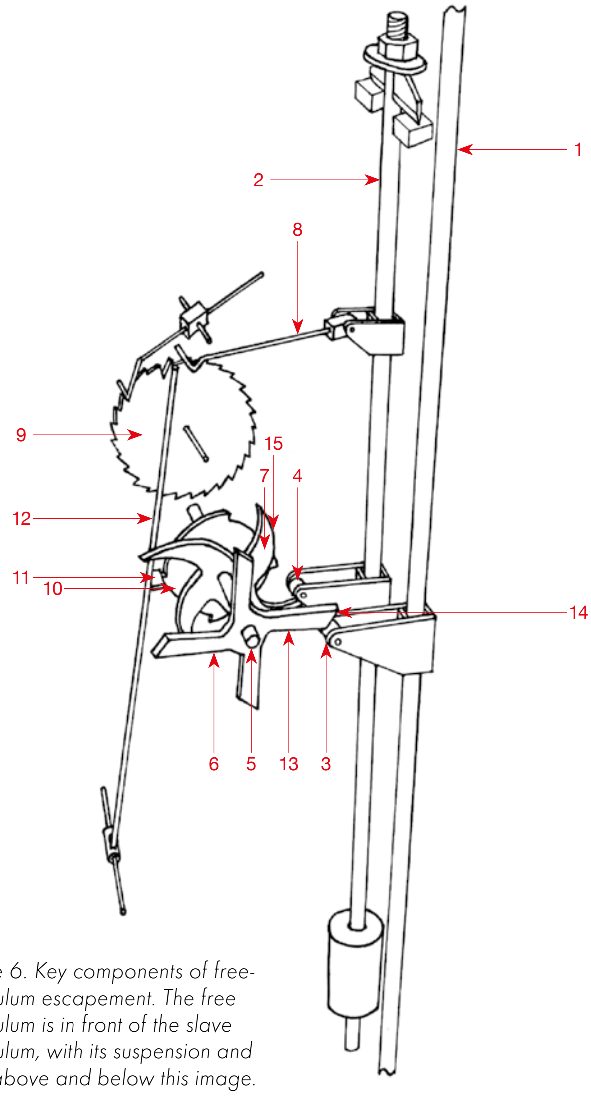

Figure 6. Key components of free-pendulum escapement. The free pendulum is in front of the slave pendulum, with its suspension and bob above and below this image.

Each lobe of the master's cam has a dead-roll surface **13** along its lower edge, and a sloped drive surface **14** at its end. The dead-roll surface is shaped as an arc of a circle with its centre at the master-pendulum's suspension point, so that it has as little effect on the master as possible. When the cam arbor is released, the master's impulse roller initially contacts the dead-roll surface, and then rolls along it until it reaches the drive surface. The drive surface applies an impulse at the centre of the master's arc, minimising any escapement error.

Each lobe of the slave pendulum's drive cam has a curved drive surface **15**. For example, this could be shaped as a part of a spiral with the cam radius proportional to the angle of rotation of the cam, so as to exert a constant force on the slave throughout the impulse. But, importantly, there is no dead-roll surface on the slave's cam.

In my clock at present, the slave has a period of $^{30}/_{29}$ seconds so that the 29-tooth count wheel rotates in 30 seconds. The master has a period of $^{30}/_{17}$ seconds.

The pendulums are suspended close to each other, so there is a concern that they might interfere with each other's oscillation. To minimise this risk, I chose the periods to be in a ratio of prime numbers, but I suspect this is overkill. It would be more convenient if the slave had a 1 second period so that the count wheel could have 30 teeth and display seconds. And the period of the master can be any integer fraction of the time between impulses, so that the master moves into the same phase relation with the slave for every impulse sequence.

Note that if the slave is allowed to swing freely, its period is a little longer than $^{30}/_{29}$ seconds. Escapement error advances the slave a little every time it is impulsed, shortening the period in which the impulse is applied, so that on average its period is $^{30}/_{29}$ seconds.

The Impulse Sequence

Recalling that a constant clockwise torque is applied to the cam arbor, when the slave's pusher moves the release lever, the cam arbor is unlatched and the dead-roll surface of the master's cam moves into contact with the master's impulse roller. At this point, the master is swinging to the right, so the roller moves along the dead-roll surface until it reaches the drive surface of the cam. Then the impulse is delivered to the master.

When the master's impulse roller moves free of the cam, the cam arbor rotates through a small clearance angle until the slave's cam contacts the slave's impulse roller. At this point, the slave is swinging to the right and so receives its impulse until, at the end of the impulse, its impulse roller moves out of contact with the cam. After a further free rotation through a small clearance angle, the cam arbor latches as the next latching surface of the latching wheel contacts the pallet on the release lever. The pendulums then swing freely until the next impulse sequence.

Synchronisation

The impulse sequence for this clock can be visualised using phase circles for the two pendulums. A phase circle represents the oscillation of a pendulum as a clockwise rotation at constant speed around a circle. The position, or phase, of the pendulum during its swing is represented by the angle at the centre of the phase circle.

The left end of the pendulum's swing is at the left-hand side of its phase circle, at 0°. The centre of its swing as it moves from left to right is then at 90°, at the top of the phase circle. The right end of its swing is at 180° at the right-hand side of the phase circle and the centre of its swing as it moves from right to left is at 270° at the bottom of the phase circle.

Figure 7 shows phase circles for the master (on the left) and the slave (on the right). The pendulums have different periods and so they rotate around their phase circles in different times; the master completes a phase circle every $^{30}/_{17}$ seconds and the slave every $^{30}/_{29}$ seconds. The pendulums, therefore, also move through different angles of their phase circles in any particular time interval.

In **Figure 7**, the impulse sequence of the clock starts on the slave's phase circle at point **A**, when the slave operates the release lever. At the same time, at point **A** on the master's phase circle, the dead-roll surface of the master's drive cam moves into contact with the master's impulse roller. The slave's phase circle shows that this happens as the slave swings from right to left, near the centre of its swing. The master's phase circle shows that at this time the master is swinging from left to right, between the left end of its swing and the centre of its swing.

After point **A**, the master continues to swing to the right, until its impulse roller moves from the dead-roll surface of its cam to the drive surface, at point **B**. The master then receives its impulse as it moves through the centre of its swing, until its impulse roller moves free of the cam at point **C**.

Looking now at the slave's phase circle, at point **C** the slave's drive cam moves into contact with the slave's impulse roller. The slave is now moving to the right, away from the left end of its swing.

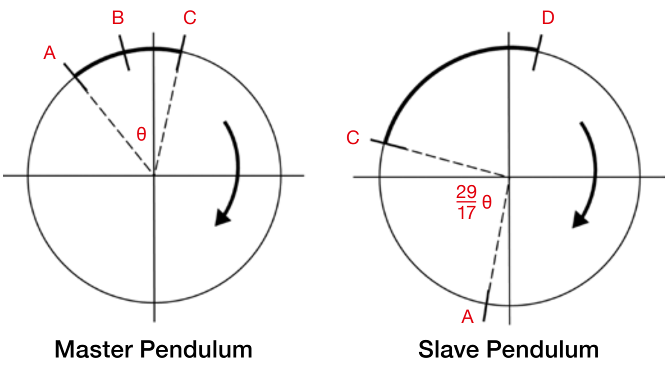

Figure 7. Phase circles for master and slave pendulums.

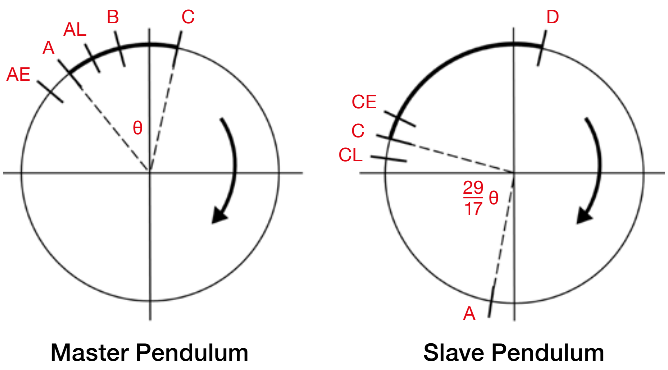

Figure 8. Phase circles illustrating pendulum synchronisation.

The time interval between points **A** and **C**, during which the master's drive cam has been in contact with its impulse roller, is the same for both pendulums, but the angular distance around each phase circle between **A** and **C** is different because of the different periods of the pendulums. These angles are marked θ and $^{29}/_{17}$ θ in the phase circles, as the ratio between the angles is 17:29.

After point **C** the slave continues to swing to the right, receiving its impulse until its cam moves out of contact with its impulse roller at point **D**. The cam arbor then rotates freely through a small clearance angle until it latches.

The phase circles in **Figure 8** illustrate in more detail how the impulse sequence synchronises the slave. In **Figure 8** points **A**, **B**, **C** and **D** are the same as in **Figure 7**.

On the slave's phase circle, points **A** and **D** (when the pusher operates the release lever and when the slave's impulse ends) are fixed by the geometry of the escapement but point **C** (when the master's impulse finishes and releases the slave's cam onto the slave's impulse roller) is not. Point **C** for the slave depends on the phases of the pendulums at that moment.

On the master's phase circle, points **B** and **C** (the start and finish of the master's impulse) are fixed by the geometry of the escapement but point **A** (when the pusher operates the release lever and releases the master's cam onto the master's impulse roller) is not. Point **A** for the master depends on the phases of the pendulums at that moment.

Suppose that points **A**, **B**, **C** and **D** represent a desired 'equilibrium' or steady-state operation of the clock. Then suppose that the slave's phase has moved behind this equilibrium position (because the slave is running a little slow) when an impulse sequence starts. The slave operates the release lever at the same fixed point **A** in its swing, but because the slave is lagging relative to the master, this point is *later* than the desired equilibrium point **A** in the master's swing. This might be at point **AL** on the master's phase circle in **Figure 8**. Therefore the master's drive cam moves into contact with its impulse roller later in its swing. But this is contact with the dead-roll surface of the master's cam and so has no practical effect on the master's oscillation.

After the master's impulse, the master's impulse roller moves free of the master's cam at point **C**. But at this point the master's impulse roller has been in contact with its cam for a shorter time than in the equilibrium situation (from **AL** to **C** rather than from **A** to **C**) because of the later release of the master's cam by the slave.

During this shorter time, the slave has moved less far than in the equilibrium situation, perhaps to point **CL**, and this is the point at which the slave's cam moves into contact with its impulse roller. The end of the slave's impulse at point D is still fixed by the slave's geometry, and so the earlier start of the impulse at **CL** increases the duration of the slave's impulse (from **CL** to **D** rather than from **C** to **D**). Because the slave's drive cam applies a force to the slave pendulum throughout the impulse, the increased impulse duration increases the escapement error applied to the slave and advances the phase of the slave towards the equilibrium point.

Suppose instead that the slave's phase has moved ahead of the master's phase when an impulse sequence starts (because the slave is running a little fast). The slave operates the release lever at point **A**, but this is earlier in the master's swing than in the equilibrium condition, perhaps at point **AE**. The master's cam is then in contact with its roller for a longer time (from **AE** to **C**), and releases the slave's cam onto its impulse roller later in the slave's swing, perhaps at point **CE**. The duration of the impulse applied by the slave's cam is therefore decreased, and the escapement error applied to the slave is decreased. This retards the phase of the slave towards the equilibrium point, relative to what it would have been in the steady-state condition.

Unlike a hit-and-miss synchroniser, the synchronisation applied by this escapement is continuously variable, and positively moves the slave's phase towards the equilibrium point whether the slave moves ahead or behind that equilibrium point. In practice, after starting the clock, it reaches steady state within perhaps 5 or 10 minutes and after that I have detected little or no variation from impulse to impulse. The feedback towards an equilibrium point seems to keep a very steady position for the initial contact point of the master's cam on its impulse roller, which is an important feature of the escapement's operation for accurate timekeeping.

As shown in the phase circles, in my clock I have adjusted the initial point of contact to be about 20° before the start of the master's impulse. Because the slave's period will vary over time differently from the master's period due to temperature or other changes (otherwise there would be no need for the master at all) the initial point of contact (**A**) of the master's cam on the dead-roll surface of its impulse roller may vary a little over time. But this is the same as in a clock with a hit-and-miss synchroniser. If an equilibrium point is initially set such that a hit-and-miss synchroniser operates on alternate impulses (hit, miss, hit, miss…), then as the slave period varies, the hit-and-miss synchroniser will require more frequent hits or more frequent misses to maintain synchronisation.

It might be suggested that in my clock the variation of the contact point of the master's cam on its impulse roller reduces the freedom of the master, but this contact point is on the dead-roll surface of the cam and so should have no effect on the master. (The Shortt and W5 clocks operate in exactly the same way. In both of these clocks, variation in the contact time between the gravity arms and the impulse arms impulsing the master pendulums, due to early or late triggering by the slave, determines the feedback through the hit-and-miss synchroniser.)

Nevertheless, as a design principle it seems desirable to minimise the length of the dead roll, and I can adjust this by varying the point A at which the slave operates the release lever, for example by adjusting the length of the pusher which actuates the release lever and/or by adjusting the slave's period.

Putting it Into Practice

Figures 9 to 16 are photographs of my clock using this free-pendulum escapement. **Figure 17** is a schematic side view of the cam arbor layout and **Figure 18** illustrates the shapes of the various cams.

Bear in mind that this clock is a prototype for developing the escapement and so please forgive mechanical short-cuts and infelicities. Having proved the principles I am designing a better-engineered version, but impatience has got the better of me and so I am writing this article first.

The escapement is held between front and rear plates (of acrylic sheet) and secured to a wooden back-board. The cam arbor is supported in ball bearings **16**, see **Figure 17**, between the plates. A clockwise torque is applied to the cam arbor by a Huygens endless loop of (black) cotton thread, wrapped around a groove **17** cut into the cam arbor. This loop forms a remontoire, passing around pulleys supporting two weights and around a large (blue) re-winding pulley. Readers may recognise the use of Meccano here.

Larger drive weights, acting on a (white) cord wrapped around a groove in the rewinding-pulley arbor, apply a clockwise torque to the re-winding pulley, **Figure 13**. The pulley is prevented from rotating clockwise by a hook at the end of a pivoted arm, mounted above the pulley, which acts on a two-lobed stop fixed to the pulley. The opposite end of the arm is connected by a string to the heavier of the remontoire weights. As the weight falls it pulls the hook from the stop and allows the pulley to rotate, re-winding the remontoire.

There are no toothed wheels in this part of the clock; the only 'gearing' is

Figure 9. Overall view of clock layout, showing larger drive weights and smaller remontoire weights.

Figure 10. Enlarged view of clock layout. Note the Z-shaped wire pusher for advancing the minute hand and the Ferguson's paradox motion work behind the dial.

Figure 11. View from the right of the escapement, showing knife-edge pendulum suspension.

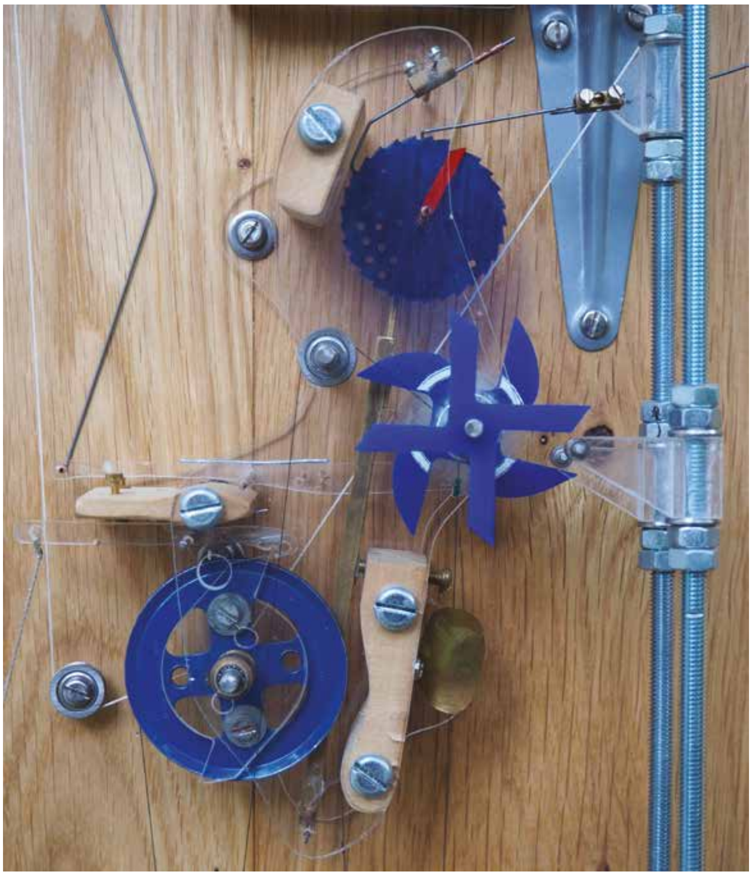

Figure 12. The escapement and remontoire. The white cord wraps around and drives the re-winding pulley arbor clockwise.

Figure 13. The escapement from the left. The white cord rewinds the remontoire and the black cotton thread, near the front of the mechanism and wrapped around the cam arbor, is the Huygens loop of the remontoire.

Figure 14. Close-up of the escapement. Pendulums free to swing with the cam arbor latched.

Figure 15. Close-up of the escapement from the left.

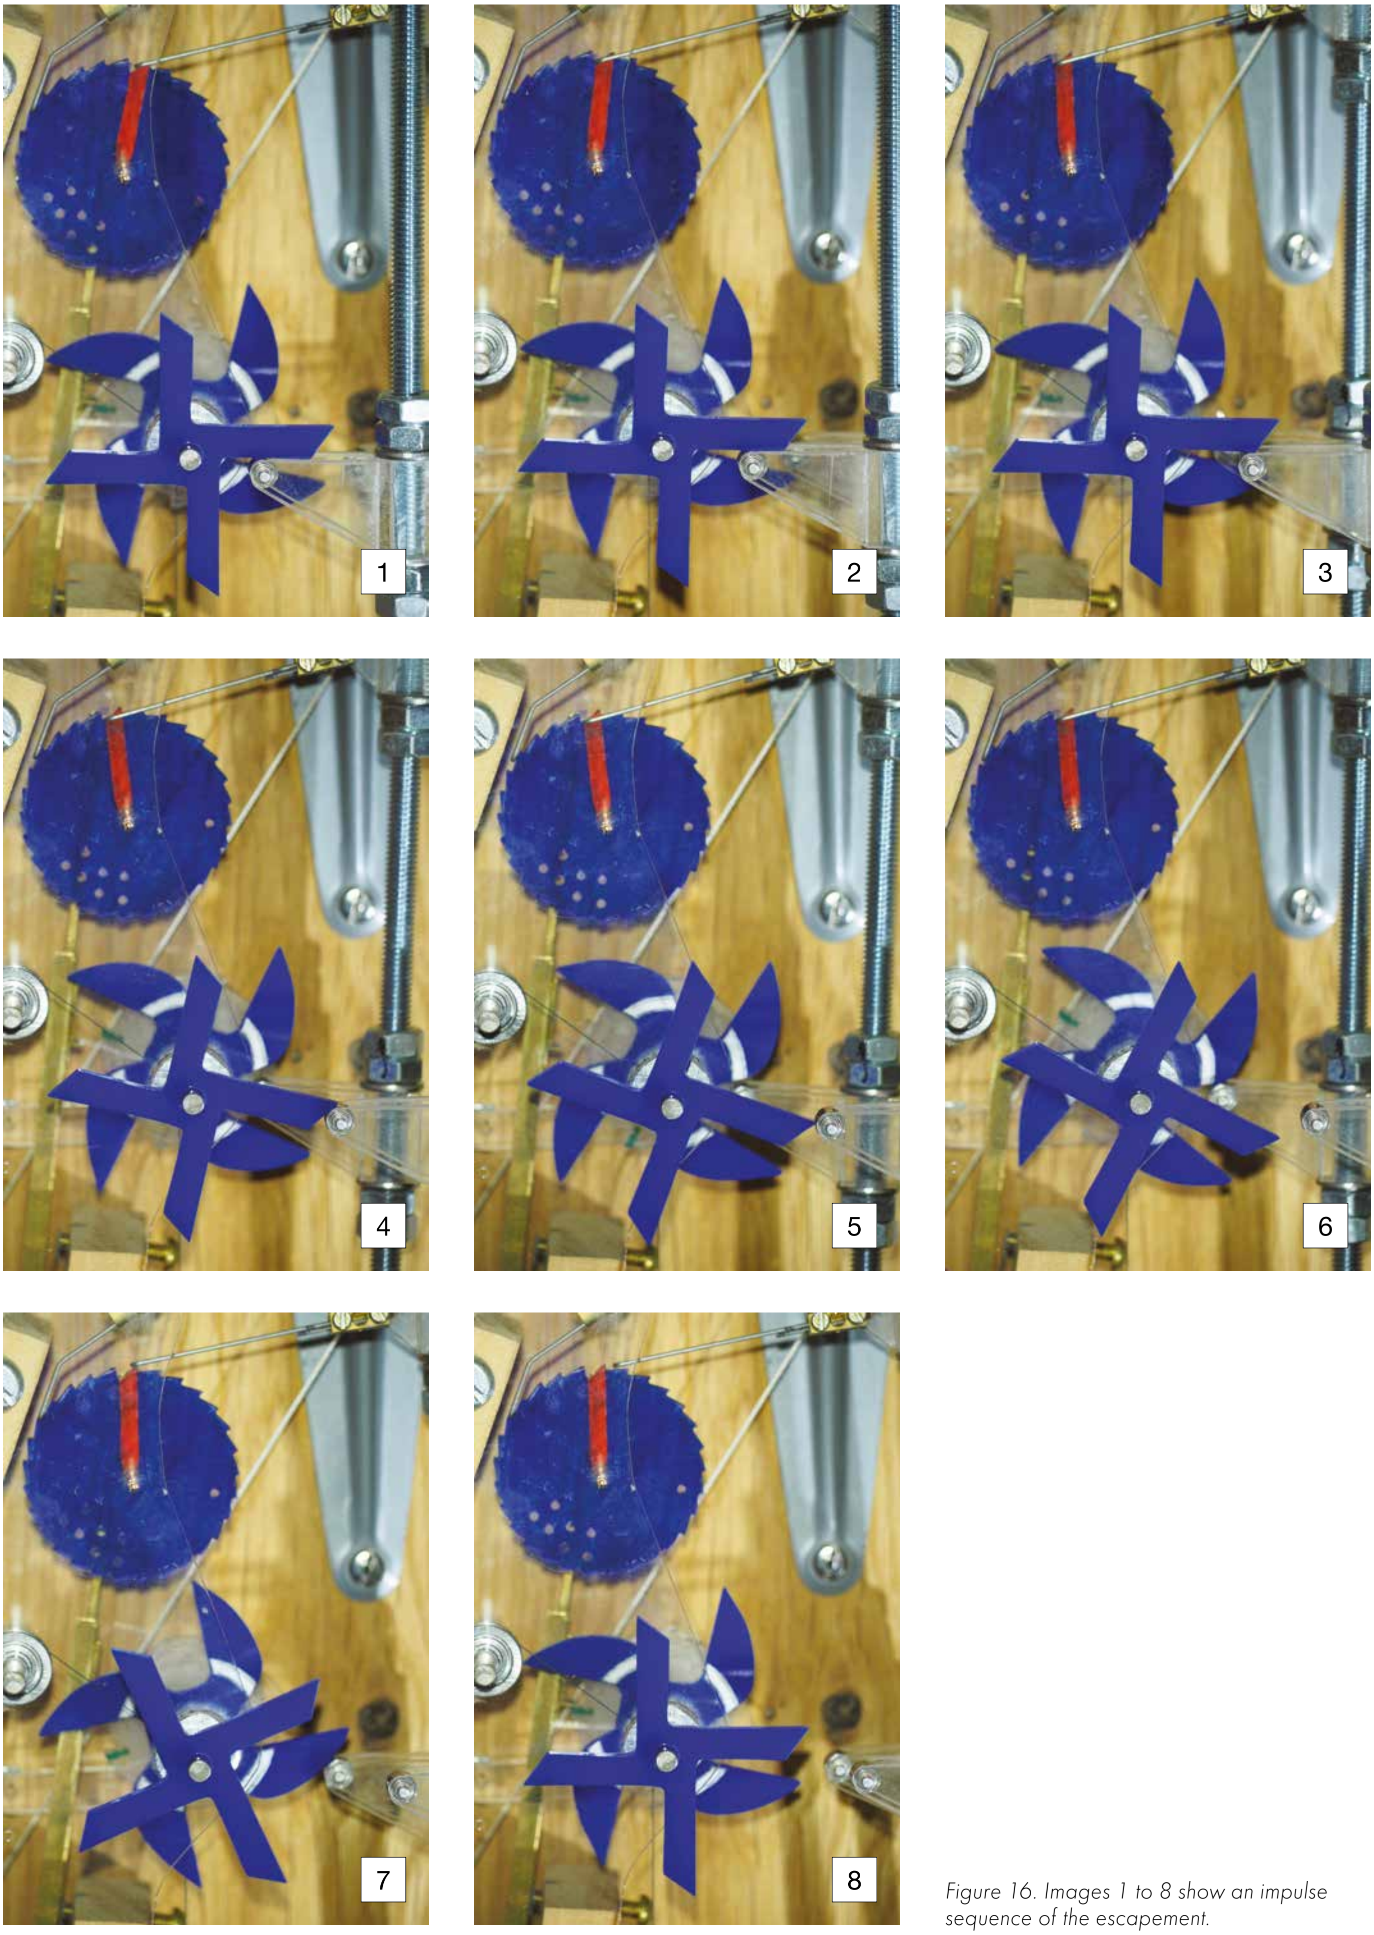

Figure 16. Images 1 to 8 show an impulse sequence of the escapement.

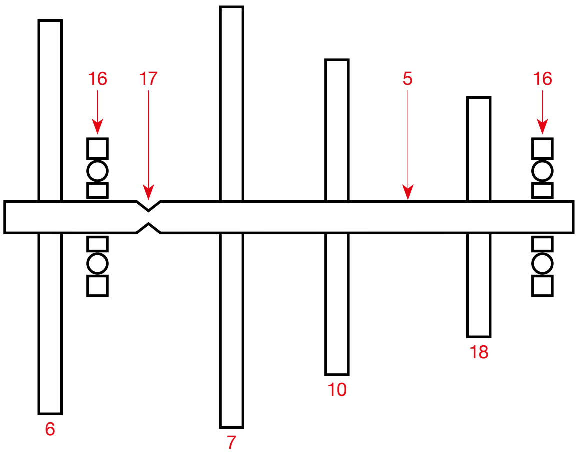

Figure 17. Schematic side view of the cam arbor. The front of the clock is at the left-hand side of this Figure.

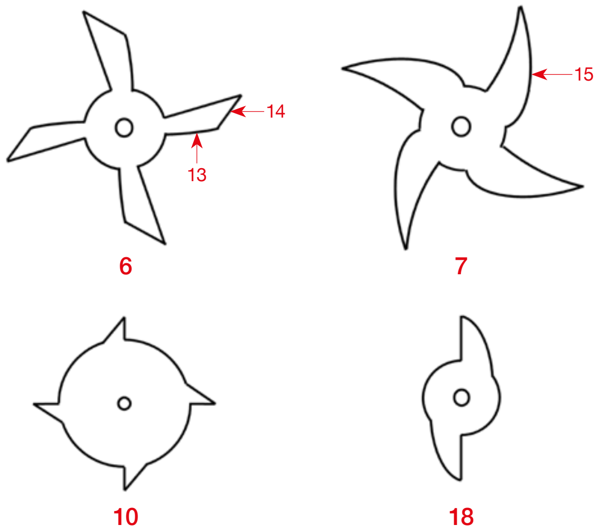

Figure 18. Schematic views of the cams and latching wheel carried by the cam arbor.

provided by the ratio between the diameter of the groove on the re-winding-pulley arbor, and the diameter of the pulley, but this is enough to allow 8-day running of the clock with a drive weight of a little less than 3 kg.

The pendulums in this clock (remember it's a prototype) are of M6 threaded steel rod, suspended on knife edges (Stanley knife blades – other brands are available). The threaded rods conveniently allow me to adjust all of the components of the pendulums to any point along their length. The knife edges are supported by wall-shelf brackets fastened to the wooden back-board. The pendulum bobs are rolls of lead sheet, seated on washers supported by locked pairs of nuts. Beneath each bob, a Nyloc nut allows fine adjustment of the pendulum period.

Each pendulum also carries its impulse roller. Each of these is a 4mm OD ball bearing supported on a steel arbor in an acrylic bracket.

The slave pendulum carries a pivoted pusher, shaped from steel wire, for advancing the count wheel.

Figure 16 shows a series of images of an impulse sequence as follows:

1. Just before the slave operates the release lever. The slave is moving to the left, and the master has just started moving to the right. Note the release lever pallet engaged with the latching wheel, and the small clearance between the master's drive cam and its impulse roller;
2. Just after the slave operates the release lever. Note that the release lever pallet has cleared the latching wheel, and the master's drive cam has contacted its impulse roller;
3. The master's impulse about to start;
4. Impulsing the master as it swings to the right through the centre of its oscillation;
5. The master's impulse roller is just moving free of its cam. Note the clearance between the slave's cam and its impulse roller at this point, ensuring that the master's impulse is complete before the slave's impulse starts;
6. The master's impulse roller has moved free of its cam, and the slave's drive cam has moved into contact with its impulse roller. The slave has passed the left-hand end of its swing and has just started moving to the right;
7. Impulsing the slave;
8. Cam arbor relatched and pendulums swinging freely until the next impulse.

An improved version of the clock will have much better-executed pendulums and supports, but the timekeeping of the clock is very encouraging even with the current arrangements.

In the side view of the cam arbor layout in **Figure 17**, the rearmost cam **18** is for driving the hands on a clock dial. This cam has two lobes which act on a cam follower during alternate impulse sequences of the clock (i.e. once a minute). The cam follower lifts a wire pusher which advances a 60-tooth count wheel behind a clock dial, attached to a minute hand. An hour hand is geared to the 60-tooth count wheel by motion work in the form of Ferguson's paradox.

The cam for operating the hands only contacts the cam follower as the cam arbor rotates during the slave impulse, so that the operation of the hands (which requires some energy from the cam arbor) cannot affect the master's impulse.

Impulse Energy

During each impulse sequence of the clock, to maintain satisfactory pendulum amplitudes, the cam arbor turns through an angle of about 11° as the master is impulsed, and through about 70° as the slave is impulsed (the rest of the 90° rotation in each impulse sequence is taken up by the clearances between the latching and impulsing steps). Therefore, maintaining the slave's oscillation seems to require a little more than 6 times as much energy as the master. Philip Woodward in *My Own Right Time* estimates that W5's slave requires 7 or 8 times more energy than the master. These figures agree very satisfyingly well.

It is instructive that so much energy is required for the slave to operate the count wheel and the release lever, and therefore how much influence an escapement can have on the oscillation of a pendulum in a conventional, non-free-pendulum clock. Avoiding the need for this energy to be delivered to the free pendulum is a key factor in the improved accuracy of a free-pendulum clock.

Design Considerations

For the best timekeeping the Q factor of the master should probably be as high as possible, but to ensure synchronisation the Q factor of the slave should perhaps be lower. For example, the lighter the slave's bob, the more effect the variation of escapement error will have on synchronisation. (This is the same as with a hit-and-miss synchroniser.) In practice, I believe that the slave should have sufficiently high Q to function well, but not so high as to resist synchronisation.

Increasing the interval between impulses may further improve the performance of the master. The energy supplied to both pendulums at each impulse would need to be greater, but the escapement error introduced into the slave would also be greater and so synchronisation should be maintained.

A problem with this escapement might be that varying the starting time of the slave's impulse (between **CE** and **CL** in **Figure 8**) to synchronise the slave also varies the total impulse energy applied to the slave. This is not a characteristic of clocks with hit-and-miss synchronisers, which in principle can adjust the phase of the slave without changing the slave's energy. In practice, I find that the variation in impulse energy required for synchronisation seems to be very much smaller than the total energy in the impulse and does not affect the performance of the escapement.

As mentioned above, the slave needs an impulse energy more than 6 times larger than the master to maintain its oscillation. Compared with this relatively large impulse energy, the variation in energy required to synchronise the slave seems to be small and to have little if any effect on the slave's amplitude.

I worried for a while that any variation of the slave's amplitude might affect synchronisation. For example, when the slave is lagging behind the master, the slave's impulse applies more energy (and escapement error) to advance the slave towards synchronisation, but this may increase the slave's amplitude and so lengthen its period due to circular deviation. These effects of escapement error and circular deviation would therefore act in opposite directions. Based on some approximate calculations, I estimate that the escapement error has about five times more effect on the slave's rate than the circular deviation, so that the escapement error strongly dominates. It was a relief that the calculations produced this result, but in practice the slave does remain synchronised effectively and so experimental observation backs up the calculated estimate.

I have put the cams for driving the master and the slave on the same arbor, which is simple, but has some consequences. One is that the same torque is applied to both cams, so that the energy delivered to each pendulum cannot be adjusted other than by altering the cam shapes. (W5 has a similar characteristic because the same gravity arm impulses both pendulums, but the Shortt clock does not.) In my clock, separate arbors could be used for the two cams, but then separate drives would have to be applied to each cam arbor and some arrangement made for the motion of the master's cam arbor to release the slave's cam arbor after the master has received its impulse.

I should perhaps comment on the choice I have made to use cams to apply impulses to rather small-diameter impulse rollers rather than, as in the Shortt and W5 free pendulums, using the interaction between an impulsing arm or a gravity arm ending in a sharp corner and a larger-diameter impulsing roller. In my clock, there is a constraint which arises from using multiple cams on a cam arbor.

When the cam arbor is latched, the slave pendulum's impulse roller must clear the rear of the slave impulse cam, and so a small-diameter slave impulse roller increases the cam-arbor rotation angle available for the slave impulse cam to drive the slave in each impulse sequence. This improves the efficiency of the clock by reducing wasted rotation of the cam arbor.

By contrast, the free-pendulum impulse roller could be increased in diameter, and the free-pendulum impulse cam could end in a sharp corner rather than in a sloped surface, but using a sloped surface applies a more constant impulse to the free pendulum over the length of the impulse. In the Shortt clock and in W5 the free-pendulum impulse is highly asymmetrical, rising from zero to an (in theory) infinite impulse force as the impulse arm falls from the impulse roller. Which is the better solution is open to discussion and perhaps needs some experiments to resolve.

Instead of using cams, I could have used gravity arms. This might be better for providing a very precise impulse to the master, but some means for resetting each gravity arm between impulses would be needed, adding complexity.

Performance

This clock is only a prototype but it runs reliably, demonstrating that the synchronisation of the slave is effective. Also, despite the sub-standard pendulum design, time-keeping to date looks good. The clock's rate seems to be stable within about a second a week, although this is over a short test period with little variation in temperature. Over a longer period the effect of temperature variation on the simple steel pendulum would certainly introduce much greater errors but the initial results are promising.

I plan to build Invar or quartz rod pendulums, at least for the master and probably also for the slave, and to acquire or build suitable timing technology to monitor performance properly. In the meantime, I look forward to any and all feedback, which may improve the next version of this clock.