Levitating Lamps (and Other Objects)

A tutorial for experimenters

By **Martin Ossmann** (Germany)

How can you make objects levitate effortlessly in the air, to all appearances? Electronicists will naturally think along the lines of (electro) magnets and mutual repulsion. Unfortunately it's far from easy to devise an arrangement that provides stable, long-term levitation. This article describes some experiments and schematics for achieving something of this kind. It also shows a method for integrated inductive energy transfer to illuminate LEDs aboard the levitating object.

Some while back the author came across a product review for a 'levitating lamp' (the Flyte Lamp [1], see **Figure 1**). This aroused both curiosity and an order for two samples, one for demonstrating and the other for 'reverse engineering'. Before these lamps arrived some time had already been spent finding out how they were likely to work [2]. The essential keyword for a Google search is 'repulsive magnetic levitation' (however off-putting this may sound, the use of 'repulsive' is correct

Figure 1. The levitating lamp going by the name of Flyte.

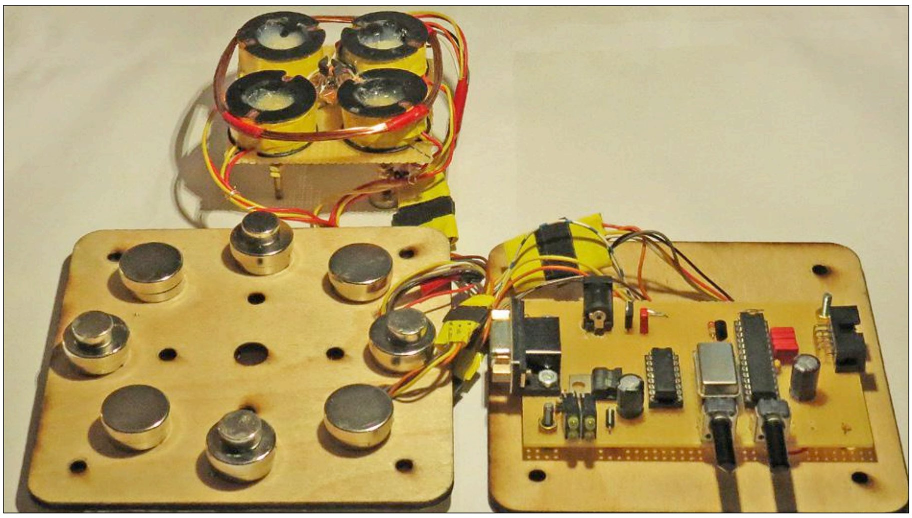

Figure 2. Components of the flying saucer.

according to the *Oxford English Dictionary*, *Ed*.). Up cropped a range of photos and videos, but alas, no 'how to' instructionals. Hence this article, in which we describe in detail how these things really work. You can read all the information needed to replicate something of this kind yourself. We'll also take a closer look at (and inside) the Flyte Lamp.

How does it all work?

Common to all implementations is the use of circular (ring-shaped) magnets to produce the opposing force against which another magnet levitates. If you get hold of a ring magnet of this kind and hold, for example, a spherical magnet centrally in the vicinity, you can feel the opposing force directly. Of course the levitating magnet will do its best to 'escape' by dodging sideways.
To prevent this from happening we use electronic regulation to stabilize the levitating object in the x-y plane. Normally four electromagnets are used, two for the x-direction and two for the y-direction. These are connected in pairs, either in anti-parallel or anti-serial. In this way we can exert force in the x- and y-directions on the levitating magnet. The position of the levitating magnet is normally registered using two or four Hall sensors.
According to one source, an Arduino is fully adequate for this regulation or control function, so the author chanced his luck using an Atmel AVR. Employing regulation is one of the ways of disproving Earnshaw's theorem [3], which states in essence that it's impossible to levitate a magnetic object using permanent magnets alone. In **Figure 2** you can see the typical components: a board with a ring of permanent magnets, a board carrying electromagnets, together with the control electronics. The complete arrangement is shown schematically in **Figure 3**.
Experience teaches that it is best to mount the permanent magnets and electromagnets on separate boards (P1 and P2), whose degree of separation *d* can be adjusted using screws. In this way you can optimize the operation of each type of magnets.

First results

If you feel like levitating some magnets yourself, you will do well to assemble a goodly collection of differing magnets (as the author did from www.supermagnete.de). When doing this, remember to factor in potential loss by breakage. If two neodymium magnets snap together with full force, they can easily self-destruct due to their brittleness. Because the electromagnets need to be placed inside the ring magnets, the ring magnet needs to be of relatively large internal diameter, which may be difficult to procure. For this reason the author used, instead of a single ring magnet, a circular arrangement of four or eight round permanent magnets. Magnets with a central mounting hole were screwed tightly onto a wooden board (Figure 2), with additional flat disk magnets placed on top of the drilled magnets for reinforcement. Using overlaid disk magnets you can fine-tune the magnetic field with ease.
The inner diameter of the ring of magnets lies between 80 and 100 mm. By these means you can achieve really substantial forces of repulsion. Once this magnetic platform has been built, we have created the first ingredient, with repulsion in the z-direction assured.

A little theory

Some of you may be wondering why we use a ring magnet, rather than a disk magnet, to produce forces of repulsion. For the answer we need to delve a little into the theory of repulsion and attraction for magnets. The magnetic field of a ring-shaped magnet is easy to calculate (using FEMM for example), as shown in **Figure 4**.
The red/green rectangles are cross-sections of the ring magnet. The axis of rotation is vertical at the center. The red/green circle symbolizes a (levitating) spherical magnet.
The problem now is that this field diagram is not very informative about the forces affecting test magnets. That's because the magnetic field itself indicates not the power force on a test dipole (= magnet) but the torsional moment. In the absence of other moments the dipole would turn as the magnetic field indicates.

The gradient (the direction and magnitude of the greatest change) is responsible for the force on a dipole, with the position of the test magnet also playing a role. With a disk magnet we now find that the test magnet always aligns itself into a position in which it is attracted in the direction of the disk. For a ring magnet things are different. There is a region in which the test magnet rotates to a position at which it is actually repelled. Put another way, the position in space (rotation about the axes) is automatically stable except for the rotation around the z-axis. Our electromagnets then take care of the x-y position and the ring magnet with the gravitational force around the z-position. The levitating object is able to rotate about the z-axis, since this degree of freedom is not fixed.

All of this can be demonstrated conveniently using an experimental setup seen in **Figure 5**. For this we use a ring magnet and a spherical magnet. The spherical magnet is free to rotate in an acrylic (Plexiglass; Perspex) tube located at the center of a ring magnet. In this way the tube restrains the sphere in the x-y direction. In the z-direction the ring magnet's forces of repulsion assure levitation, with the spherical magnet rotating in a position where the repulsion is at its maximum.
Incidentally it's interesting to note what happens when you place the setup in Figure 5 on your head. You would expect the force of repulsion to expel the sphere from the acrylic

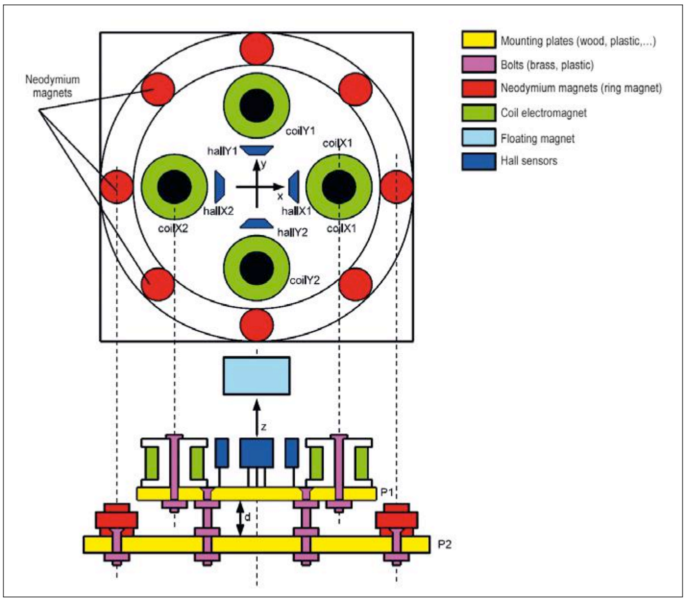

Figure 3. Schematic diagram showing mechanical construction.

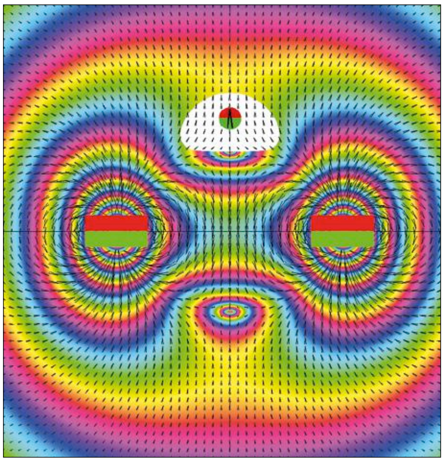

Figure 4. Magnetic field of the ring magnets.

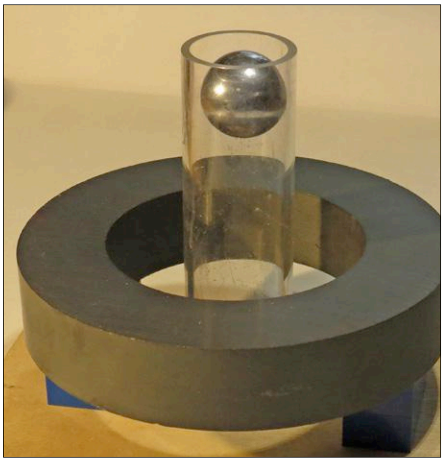

Figure 5. The sphere magnet levitates above the ring magnet.

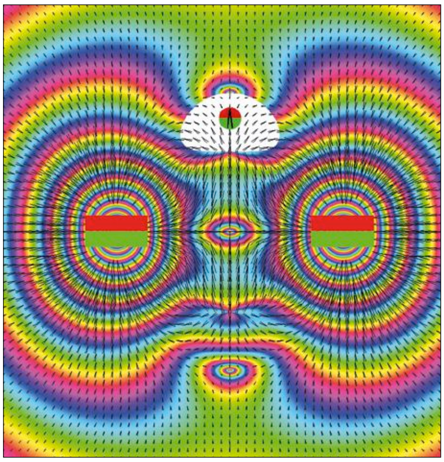

Figure 6. Energy field of the ring magnet.

tube and drop to the ground. But instead the sphere moves only a small distance away from the ring magnet, remaining in the tube, because at this greater distance it is now attracted by the ring magnet. The sphere is then located just above the white region seen in Figure 4.

If you replace the ring magnet with a disk magnet, the spherical magnet will always slam into the disk magnet with maximum impact, because it will rotate in space until the attraction is at maximum. The fact already made plausible can also be illustrated with a 'pseudo field diagram' (**Figure 6**). This shows the force that arises when the test dipole has rotated into a position in which no further moment of rotation affects it. In Figure 6 you can see clearly how at the center, above the ring, a not particularly large region of repulsion (indicated in white). This small area is where our levitating magnet will linger later on.

It is these relationships that mainly determine the height at which the levitating magnet hovers above the ring magnet. Once you have built the magnetic board, you can already, for example, use a spherical magnet to get a pretty good feel for how high the sphere will float later. The size of the white region is not dependent on the **strength** of the ring magnet but upon its **diameter**. To increase the hovering height, you must use a larger ring of magnets. With a stronger magnet, the height can be increased only within the white range.

X-Y magnets

Our next investigation with the electromagnets is into x-y stabilization. The author used either four or eight electromagnets, wound using 0.25 mm enameled copper wire. For the winding formers you can use either a P36/22 coil bobbin (1200 turns) or a homebrew (3D-printed) coil former with components as in **Figure 7** and **Figure 8** (1600 turns). These coils are usable with operating voltages up to 12 to 15 V. The four coils are mounted on a separate board, made from epoxy PCB or perfboard material for instance.

Hall elements

To monitor the position of the levitating magnet we use Hall devices of type SS496A or SS495A. You can use two or four sensors, mounted centrally either horizontally or vertically. The author obtained best results with four Hall sensors mounted vertically, with the software polling two oppositely-located sensors in a bridge arrangement so to speak.

The sensors are placed on the upper surface of the electromagnets. The Hall elements are wired as shown in **Figure 9** and

Figure 7. Winding formers made using a 3D printer.

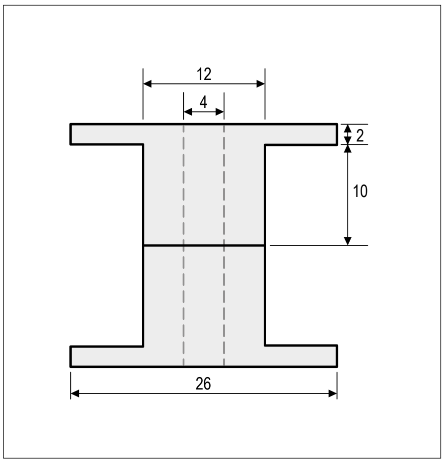

Figure 8. Dimensions of the winding former.

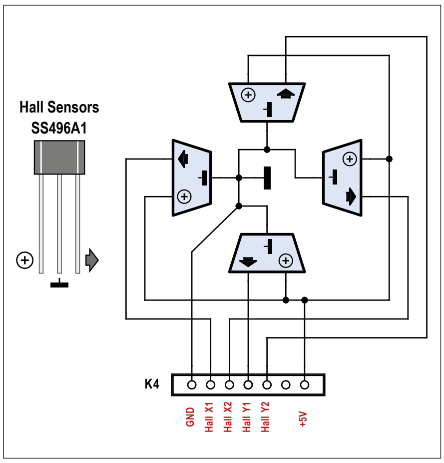

Figure 9. Wire connections to the Hall sensors.

linked to the control electronics using the indicated (6-way) plug connector. Once you have assembled the electromagnets together with the sensors, you can make the first tests using the control circuitry that follows.

Control electronics

The entire electronic circuitry is shown in **Figure 10**. An ATmega microcontroller generates a pair of PWM signals that control the L293D motor driver IC3. This dual full-bridge driver powers the two pairs of coils. Using the LEDS wired in anti-parallel, you can check the direction of current flow in the coils at any given time. If you hold the levitating magnet firmly and move it using the Hall elements, then which LED lights up will depend on the polarity of the magnets.

Four A-D inputs are used to evaluate the Hall elements, two A-D inputs digitize the voltages of the two pots R1 and R2 that are used for adjusting the two constants K_p (proportional component) and K_d (differential component). Transistor T1 and capacitor C13 provide wirefree energy transfer. Their function is explained in detail later.

In the software we have two independent PD (proportional-differential) controllers, one each for the x- and y-directions. With some simple RS-232 interface circuitry you can keep an eye on how the control software operates using RS-232 at 115,200 bit/s. The following output values are produced: control constants K_p and K_d, positions x and y, manipulating or correcting variable x and y. For readers who enjoy experimenting, some software for the ATmega is available as a free download [4].

By shifting the levitating magnet manually you can use this protocol output to check whether the sensors are functioning correctly and whether the controller produces the correct positioning signals when the magnet is displaced from the center.

In the first phase of testing it was not yet clear how much current the electromagnets would require. For this reason, instead

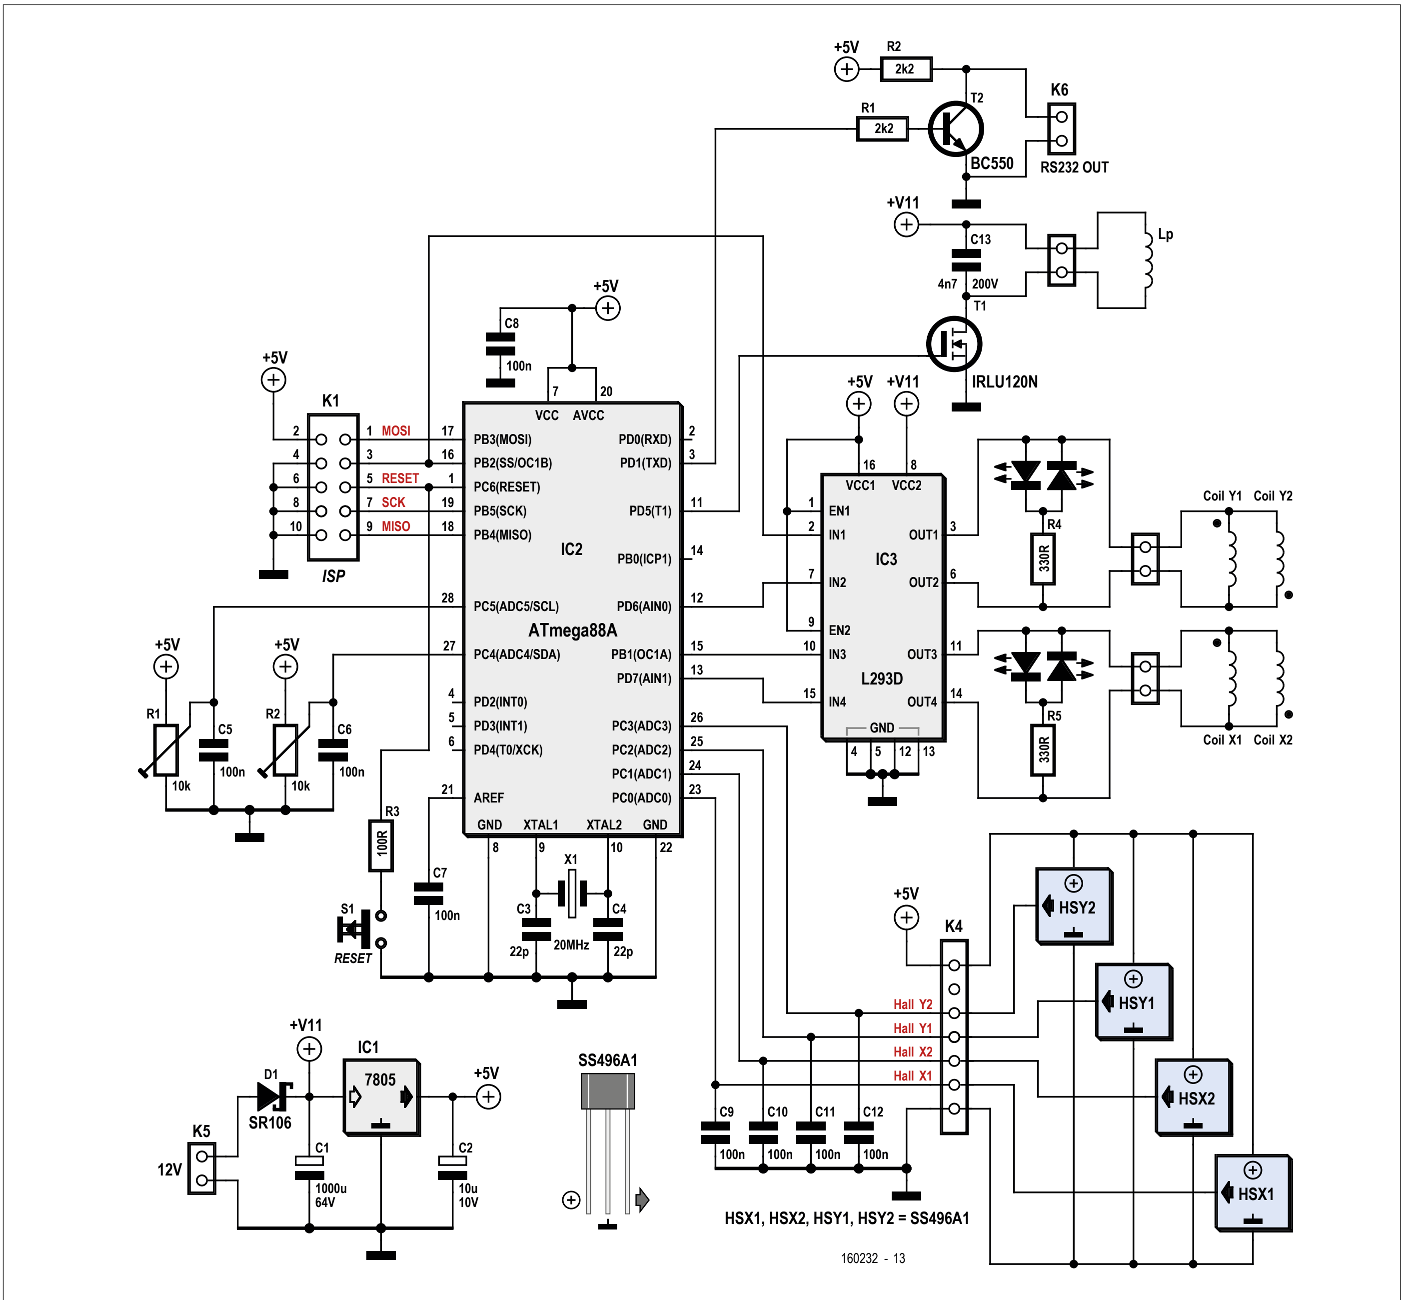

Figure 10. Schematic for the control electronics.

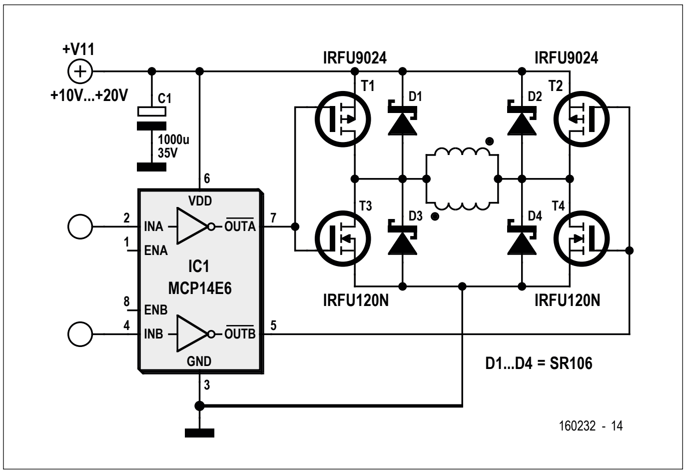

Figure 11. Output booster using MOSFETs.

Figure 12. Electromagnet board for the flying saucer.

of using the L293D as bridge driver, two discrete full bridges were used as per the schematic in **Figure 11**.

With this circuit you can definitely push several amps through the coils. For testing at higher currents you will then do better to replace the L293 with this circuit (times two of course).

Testing the electromagnet board

The first test is carried out without the permanent magnet board, so that only the electromagnets produce a field (**Figure 12**). First we test the *x*-control circuit alone, to do which we link only the x-coils and all four sensors to the control PCB. Potentiometer K_p is turned to halfway round but pot K_d remains at its left-hand end stop position (ground). If you now cautiously move a disk magnet of, say, 20 mm diameter centrally over the magnets in the *x*-direction, you can then detect and feel the operation of the regulation circuitry.

If the coils are connected the right way round, then except at the center position you will notice a weak force of repulsion that tries to move the magnet back into the center position (if the magnet is at the correct height). With the wrong polarity the control function will draw the magnet away from the central position, which of course is precisely what we don't want. In the same way you can now take care of getting the correct polarity for hooking up the *y*-coils. Once both coils are connected, you can note how the control circuit holds the magnet centrally.

Some designers mount the Hall sensors flat on top of the coils as seen in **Figure 13**. However, the author did not get good results like this because the strong magnetic field in the z-direction drove the sensors almost into saturation.

Construction and commissioning

In the assembly process the clearance between the magnet board and electromagnet board (measurement *d* in Figure 3) plays a decisive role. The distance must be adjusted in such a way that the levitating magnet hovers in the z-axis precisely where the electromagnets provide the best *x*-*y* axis control. If needs be, you must try out several different distances repeatedly. An error will be evident when the regulation fails to take control properly.

Now comes the actual process of putting our levitator into service, by setting the regulation parameters K_p and K_d. We start with K_p=0 (left-hand end stop) and K_d=0 (left-hand end stop). Next we hold the levitating magnet centrally above the electromagnets, where it will shortly float and hover. This process fixes the *x*- and *y*-directions to some degree, but not entirely.

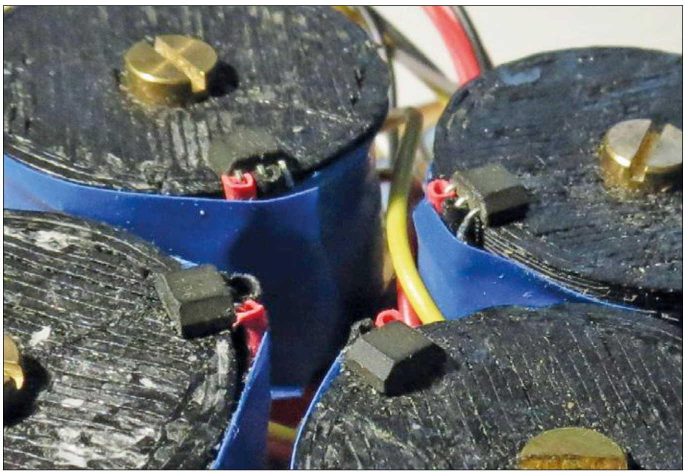

Figure 13. Low-profile Hall sensors mounted horizontally.

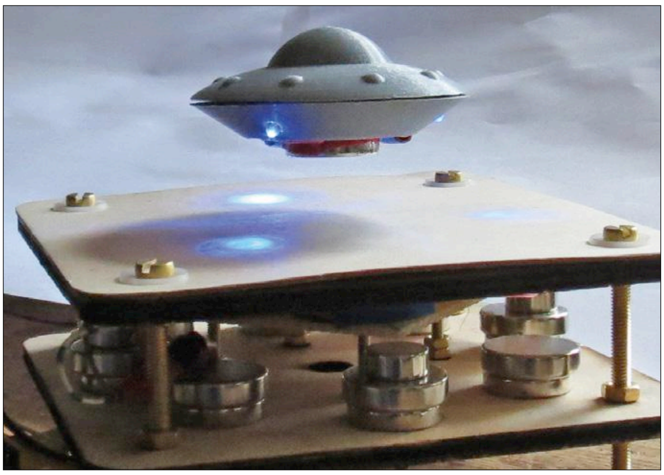

Figure 14. Utterly cute: a real-life flying saucer.

Then we turn the K_p pot slowly. You will notice how the regulation cuts in and tries to restrain the magnet centrally.
K_p is now advanced until this stabilizing effect is felt clearly. If you let go of the levitating magnet, however, the magnet will begin to oscillate (vibrate) and then slam down towards the ring magnet. To prevent this, we rely on the differential component of the controller.

So we adjust K_d just far enough to eliminate the oscillation. If this is unsuccessful, you need to readjust K_p or even the distance *d* between the magnet and electromagnet boards. By their nature, your very first attempts will be the most difficult but over time you will develop some experience.
The proportional-differential regulation algorithm is shown in **Listing 1**. You can calculate using integer figures throughout; a sampling rate of 1 kHz is then possible.

Using a 3D printer the author created a flying saucer, which even incorporates three illuminated LEDs as 'headlamps'. In **Figure 14** you can see the levitating UFO, which hovers at an elevation of about 15 mm above the electromagnet.

Increasing the clearance

To raise the amount of elevation a bit we built another example, increasing the diameter of the ring of permanent magnets, with eight electromagnets and eight permanent magnets (**Figure 15**). This made it possible to increase the vertical clearance to at least 30 mm or so.

Inductive energy transfer

To transfer energy inductively we frequently employ concepts in which both the primary and the secondary windings of transformers are brought into resonance using appropriate capacitors. Voltage and current forms are then of sinewave format. For this reason the author was mighty surprised when he measured the voltage on the primary winding of the Flyte Lamp (which by then had arrived) with the aid of an auxiliary transformer. It looked like the representation in **Figure 16**, which showed juxtaposed half-sinewaves. The author had to rummage deep in his memory before he had a light-bulb-moment hunch that this could be a Class E converter [5].
In power electronics a Class E converter is often put to work in applications involving very high frequencies. Schematically it looks like the circuit in **Figure 17**. Coil L1 is extremely large and produces an approximately constant current flow. L_r and C_r form a resonant network that connects load *R* to the circuit. If you wish to provide a DC load, then you need to add a bridge rectifier.

The capacitor C_p connected in parallel with switch *S* is highly unusual. Normally a capacitor of this kind would cause colossal interference in power electronics, as it is short-circuited when *S* is closed and its stored energy is wasted entirely.
In the Class E converter we make a virtue out of necessity. Let's start the circuit description when switch *S* is passing current and capacitor C_p is not charged. In the usual interpretation of this circuit a positive current (coming from L1) then flows through the switch. Now we open switch *S*. As the capacitor is initially discharged, the voltage at the switch rises only gradually, creating a zero voltage switching situation.
The resonant network ensures that a kind of half-sinewave

Listing 1. The proportional-differential control algorithm.

```
// Calculate errors in the x-direction:
errorX=refValueX-HallX;
// Difference between new errors and previous
   errors:
dErrorX=errorX-prevErrX;
// Calculate control value:
ctrlX=(errorX*Kp+dErrorX*Kd)/4096 ;
// Output to coil:
setCoilX(ctrlX);
// New previous value = new value:
prevErrX=errorX;
```

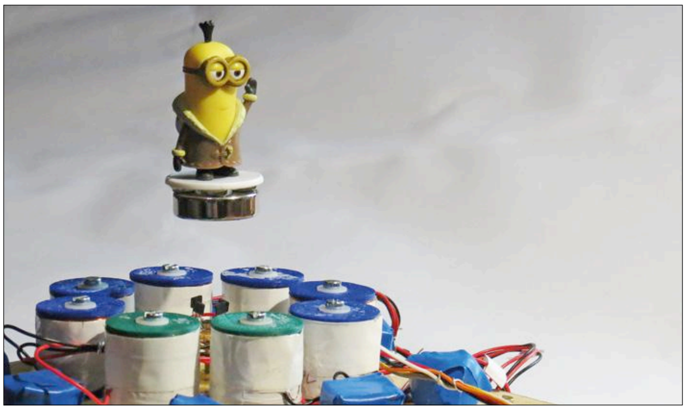

Figure 15. Here's how we achieve an elevation clearance of 30 mm.

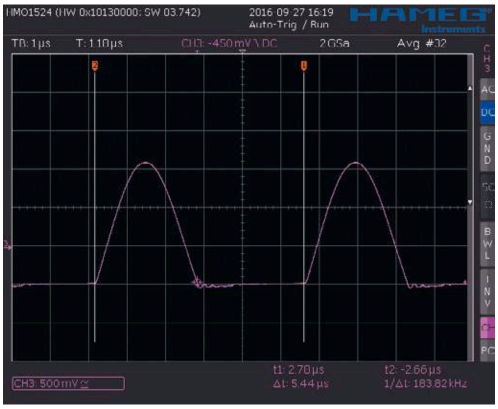

Figure 16. Induced voltage with the Flyte Lamp.

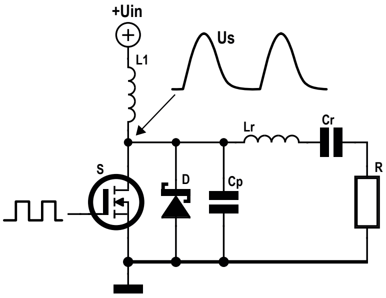

Figure 17. Class E converter shown schematically.

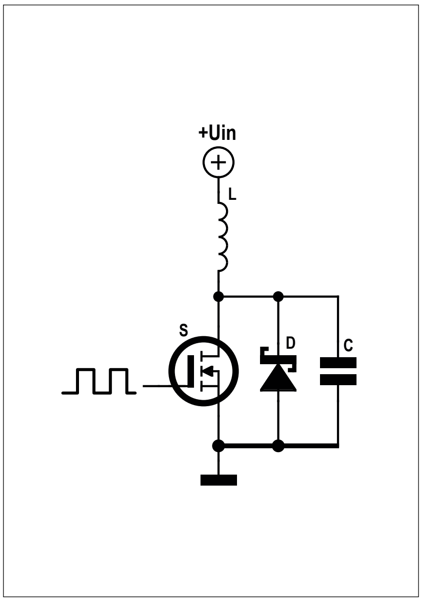

Figure 18. Simplified Class E stage.

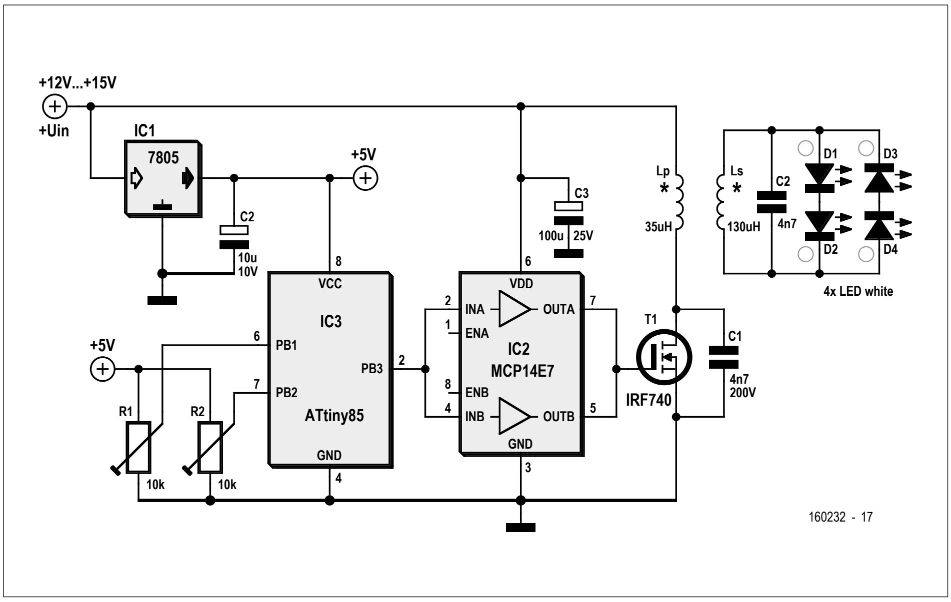

Figure 19. Test circuit of the Class E converter.

appears as the voltage form on switch *S*. When the voltage then goes negative, diode D becomes conductive and the voltage at capacitor C_p is zero once more. This holds good for a certain time, during which switch *S* is closed again. We also get zero voltage switching at switch-on time and the whole sequences starts again from the beginning. Further details can be found in the literature [5].

You can now slim down the Class E converter and construct it as shown in **Figure 18**. The combination of coil *L* and capacitor *C* makes up the resonant structure that is excited by the transistor *S* to create oscillations of the form shown in **Figure 16**. The voltage at the drain of the switching transistor creates a multiple of the supply voltage in the process. Incidentally you will find the same circuit used in the inductive hobs of electric cookers.

For test purposes a circuit along the lines of **Figure 19** was developed. The ATtiny is used here as a flexible PWM generator. The PWM signal drives switching transistor T1 by means of a gate driver. Primary and secondary windings have the values indicated (L_p = 15 turns of 0.4 mm enameled copper wire, 80 mm diameter, L_s = 50 turns of 0.2 mm / 32 AWG enameled copper wire, 30 mm diameter). The coupling factor *K* was determined by measurement as approximately 0.08.

Figure 20 shows the test setup of Figure 19, which is able to light four white LEDs spaced about 20 mm apart without difficulty. The same dimensioning was then applied to the flying saucer. Because the Class E converter achieves zero voltage switching at power-on and power-off, the switch turns on and off relatively slowly. For this reason you can also manage without the gate driver and let the microcontroller trigger a logic level MOSFET directly. In Figure 10 L_p and C13 from our microcontroller circuitry form the resonant circuit switched by T1. For the diode we employ the body diode inside the MOSFET. The frequency is about 250 kHz.

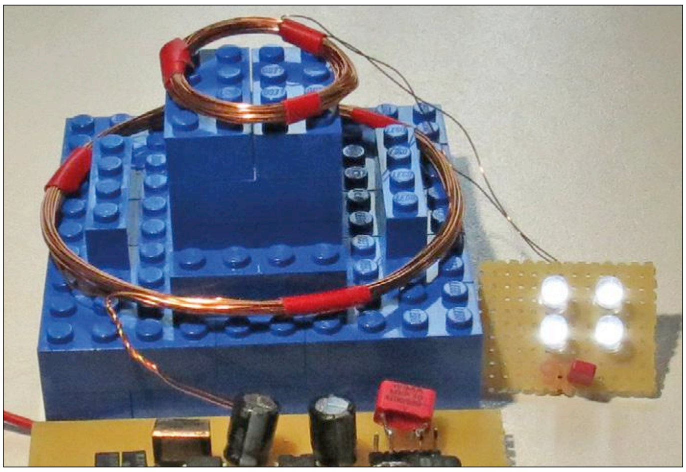

Figure 20. Testing inductive energy transfer.

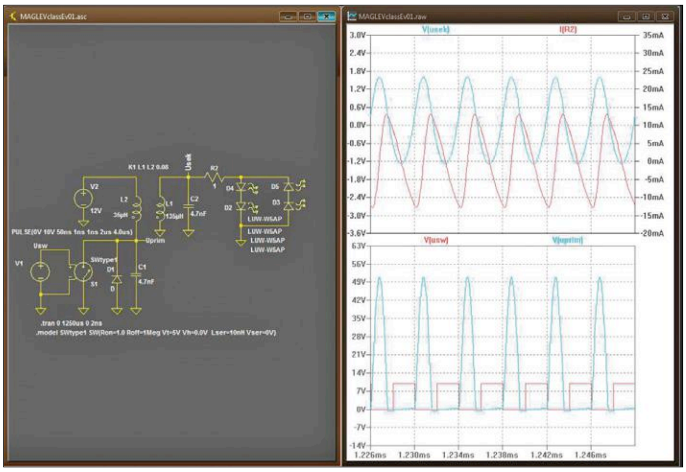

Figure 21. LTspice simulation.

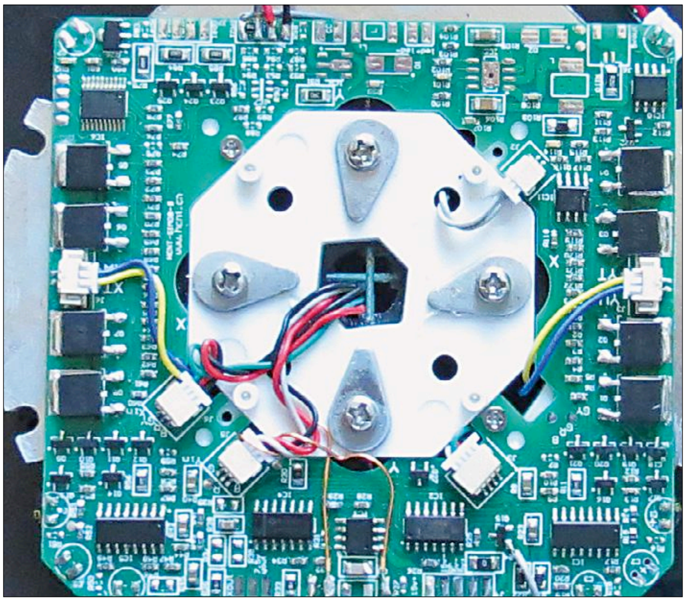

Figure 22. SMD side of the Flyte PCB.

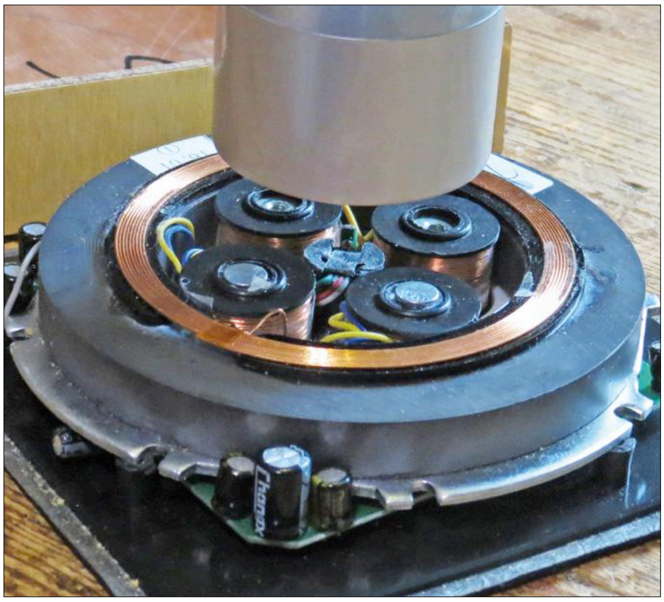

Figure 23. Flyte electronics seen from above, showing the electromagnets.

Before the converter was built its behavior was simulated using LTspice. The circuit and curve shapes arising are shown in **Figure 21**.

The Flyte Lamp analyzed

The lamp is easy to disassemble by unscrewing four screws. The lower side of the PCB and its array of SMD components are seen in **Figure 22**. The two times four 'fat' transistors make up the two full bridges used for controlling the two pairs of magnet coils. The IC labeled XKT-510 is evidently responsible for wirefree energy transfer [6]. The arrangement of the transmitting coil suggests use of the Class E converter discussed. There are also a few unidentifiable ICs on the board plus a microcontroller made by ST. The top of the board is shown in **Figure 23**.

The PCB contains the ring magnet as well as the four electromagnets. These have an iron core that appears to be adjustable (presumably for optimizing the position of the levitation magnet). In the center there are probably two Hall elements for positioning purposes. Clearly visible is the large diameter air-coil with relatively few turns, which handles the inductive energy transfer. The entire circuit is powered with 15 V. In point of fact this analysis of the Flyte Lamp reveals no great novelty that is not in our tutorial for constructing a flying saucer yourself. A patent covering another arrangement of this kind can be found in [2].

By the way, the author looked in vain for a CE conformity symbol on the device (the CE mark indicates conformance with relevant European product legislation, similar to the FCC declaration of conformity in the USA). It is even questionable whether a Class E converter constructed like this could fulfill relevant EMC requirements. There are no EMC filter components on the PCB in any case. ⏮

(160232)

Web Links

[1] **Flyte Lamp**: https://flyte.se

[2] **Eberhard Waffenschmidt, Peter Visser, Patent: Floating lamp, WO 2010150207 A1**: www.google.com/patents/WO2010150207A1?cl=de

[3] **Earnshaw, S., On the nature of the molecular forces which regulate the constitution of the luminferous ether., Trans. Camb. Phil. Soc., 7, pp 97-112 (1842)**: www.mit.edu/~kardar/research/seminars/Casimir2010/pdf/EarnshawPaper.pdf

[4] **Software for the ATmega**: www.elektormagazine.com/160232

[5] **A high-efficiency Class E inverter – computer model, laboratory measurements and SPICE simulation, Z. Kaczmarczyk**: http://bulletin.pan.pl/(55-4)411.pdf

[6] **XKT-510 IC**: http://elecfreaks.com/store/download/XKT-510.pdf

[7] **UFO video**: https://youtu.be/RE9qDVLKxS4