RepRap – the replicating rapid prototyper

Rhys Jones†, Patrick Haufe‡, Edward Sells§, Pejman Iravani†,
Vik Olliver¶, Chris Palmer‖ and Adrian Bowyer†,*

†*Department of Mechanical Engineering, Bath University, Bath BA2 7AY, UK*
‡*Faculty 5, Department of Biomimetics/Biological Materials, Hochschule Bremen, University of Applied Sciences,
Neustadtswall 30, D-28199 Bremen, Germany*
§*Bits From Bytes Ltd., Unit 17 Hither Green Industrial Estate, Clevedon, North Somerset BS21 6XU, UK*
¶*Diamond Age Solutions Ltd., 72 Warner Park Avenue, Laingholm, Waitakere, New Zealand*
‖*Boatmans Row, Astley Green, Tyldesley, Manchester M29 7JQ, UK*

SUMMARY
This paper presents the results to date of the RepRap project – an ongoing project that has made and distributed freely a replicating rapid prototyper. We give the background reasoning that led to the invention of the machine, the selection of the processes that we and others have used to implement it, the designs of key parts of the machine and how these have evolved from their initial concepts and experiments, and estimates of the machine's reproductive success out in the world up to the time of writing (about 4500 machines in two and a half years).

KEYWORDS: Self-replicating machine; Rapid prototyping; Additive fabrication; Biomimetics; Mutualism; Open-source; Free software; Fused-filament fabrication.

1. Introduction
RepRap is an open-source self-replicating rapid prototyping machine. It is a robot that uses the fused-filament fabrication** to make engineering components and other products from a variety of thermoplastic polymers. RepRap has been designed to be able to print out a significant fraction of its own parts automatically. All its remaining parts have been selected to be standard engineering materials and components available cheaply worldwide. As the machine is free†† and open-source, anyone may – without royalty payments – make any number of copies of it ether for themselves or for others, using the RepRap machines themselves to reproduce those copies.

In this paper, we briefly examine the terminology and history of artificial reproduction, and then describe the biomimetic genesis of the RepRap machine, its original design, how and why that design has changed into its current form, RepRap's global adoption and use, and the commercial offshoots and spinouts from it.

1.1. Terminology
Historically, the terminology used in the field of self-reproducing machines has sometimes been unclear, with different meanings being ascribed to the same terms. In an attempt to bring some systematisation to this, we will define key terms for use in this paper at least.

Kinematic machine: A physical machine that is composed of fixed and moveable parts. This term makes a distinction between real machines and software models (which are frequently used for simulation). In what follows, we take the words "kinematic machine" to include living organisms.

Self-replication: We start with the idea that *self-replication* could mean an imaginary Platonic process by which a kinematic machine was able to create an exact copy of itself. The Second Law of Thermodynamics and Shannon's theorem[1] show that information cannot be copied without loss or error indefinitely, implying that the idea of an exact replicator is impossibile. (Of course, it is the errors that drive Darwinian evolution.) Whilst it is philosophically and poetically useful to have words for impossible ideas, here we reduce the strength of the word "replication" to give it an engineering meaning: *a copy within specified tolerances that will work as well as the original*.

Self-reproduction: A process by which a kinematic machine is able to create an approximate copy of itself, perhaps with either insignificant or significant errors. All living organisms are self-reproducers. The specified-tolerances and works-as-well distinctions between *replication* and *reproduction* follows through the definitions below, and the rest of this paper. Replicators are a subset of reproducers.

Self-manufacturing: The ability of a kinematic machine to make some or all of its own parts from raw materials. This clearly prompts a requirement for a definition of "raw": Is an etched printed circuit board a raw material? Or a uniform copper-clad board? Or some copper, some glass, and some epoxy resin? Forensically, many Gordian knots of this sort

---

**Fused-filament Fabrication (FFF) is sometimes called Fused Deposition Modelling (FDM). But this latter phrase is trademarked by Stratasys Inc., and so is not in unconstrained use. FFF was coined by the members of the RepRap project to give a synonymous term that could be used by all unrestrictedly.

††In this paper, we always use the word "free" to convey both the meanings that it has in the free software discourse: "Free as in freedom", and "Free as in beer".

Fig. 1. Schematic of von Neumann's kinematic reproducer from ref. [4].

are cut by asking, "Would a reasonable person say it is so?" and leaving it at that. We adopt the same approach as the law.

Self-assembly: This refers to the ability of a kinematic machine to manipulate a series of parts into an assembled copy of itself.

Autotrophic self-reproduction or self-replication: The ability of a system to make a direct copy of itself from raw materials without assistance. As yet, no artificial autotrophic self-reproducing kinematic machine has been made. However, examples exist in biology (see Section 2). For a kinematic machine to achieve autotrophic self-reproduction, it must contain a number of critical subsystems. One attempt to identify these subsystems was undertaken in Freitas and Merkle's "Map of the Kinematic Replicator Design Space" in their comprehensive book,[2] which identified 137 design properties in order for autotrophic self-reproduction to be possible.

Assisted self-reproduction or self-replication: A kinematic machine that includes at least one but not all of the critical subsystems required for autotrophic self-reproduction or replication and so needs human (or other) intervention to reproduce.

By these definitions, RepRap is a kinematic assisted self-replicating and self-manufacturing machine, as we shall show below.

1.2. Artificial reproduction

The concept of self-reproducing kinematic machines has intrigued some of the world's greatest minds for generations. However, the first person to formalise thoughts on the subject was von Neumann in the middle of the last century.[3] Much of von Neumann's work concentrated on his cellular machine, a theoretical and mathematical model, and records of his research into a kinematic (physical) self-reproducing machine are scarce and often informal. Much of the outline presented here is based on the summary in the review by Freitas and Merkle.[2]

Von Neumann's kinematic reproducer, as illustrated by Cairns-Smith[4] in Fig. 1, consists of five distinct components, namely a chassis (c), a set of instructions (I), some form of machinery (m), a controller (r) and, finally, a sequencer (s).

In order for the kinematic reproducer to function properly, it is required that it resides in a stockroom containing an unlimited quantity of spare parts. The kinematic machine features a mechanical appendage that is able to gather parts at random from this stockroom; the randomly selected part is inspected and compared with the kinematic machine's instructions. In the event that the part is not required, it is replaced in the stock room and the process is repeated until a required part is found. This process is then repeated to find the next required part, and the two parts are connected together using the mechanical appendage. This cycle continues until a physical copy of the kinematic machine is produced; at which point, the instructions are copied in to the memory in the child kinematic machine before it is finally activated.

In devising his kinematic reproducer, von Neumann ignored any fuel and energy requirements. Even so, with a part-count of the chassis estimated at 32,000, the feasibility of the device was poor. Nevertheless, the concept did at least demonstrate the principle of a self-reproducer, and has inspired many people to do further research. Most of this work may be broken down into three distinct subsets. Using the above-given definitions, these subsets are as follows:

1.2.1. Autotrophic self-reproducers. With limited success so far in the areas of self-assembly and self-manufacturing, an artificial autotrophic self-reproducer remains an unachieved utopia for the subject. Whilst theoretical work has been undertaken in the area, all the concepts presented so far are extremely vague on the *engineering* involved in artificial reproduction, being described by Dyson thus: "We don't have the science yet; we don't have the technology".[5]

1.2.2. Self-assembling kinematic machines. Some of the most elegant work into self-assembling kinematic machines using special pre-made parts was conducted by Roger and Lionel Penrose in designing their so-called *block reproducers*.[6] Perhaps the biggest achievements of their design are its neatness and simplicity.

The block reproducer (Fig. 2) consists of a series of wooden blocks that are placed on an agitating surface. The design of the blocks is such that an interlocking profile exists on each block. "Brownian-motion" is induced into the parts by agitating the surface, enabling the locking profile to be utilised for completing the assembly process. The

Fig. 2. A 1-D self-reproducing kinematic machine made from parts of two kinds from ref. [6].

Penroses also designed a more complicated two-dimensional reproducing kinematic machine along similar lines.

Further work into self-assembling processes was conducted by Moses,[7] who developed a self-assembling kinematic machine in the form of a Cartesian manipulator based on 16 types of snap-fit parts. Similarly to von Neumann's kinematic reproducer, it was able to build a copy of itself if supplied with sufficient parts. However, whilst the concept proved promising, the structure of the design lacked stiffness, leading the machine to need external assistance to complete its reproduction cycle. But, inspired by this success, the world's first semi-autonomous, limited-part, self-assembling kinematic machine was created in 2003 by Suthakorn *et al.*,[8] with an assembly time of just 135 s. It consisted of an original robot, subsystems of three assembly stations and a set of subsystems from which replicas of the original robot were assembled. In 2005, Zykov *et al.*[9] made a system consisting of cubes split along a diagonal where each half-cube could rotate relative to the other in that split plane. The cube's faces were fitted with electromagnets. Stacks and other arrangements of these could be made to reproduce themselves if fed with a supply of similar active cubes, with the stack acting as a robot arm when the split faces were rotated.

1.2.3. Self-manufacturing kinematic machines. To date, the amount of work on self-manufacturing process has been limited. The main research in this area was conducted in 1980–1982 by the Replicating Systems Concepts Team at NASA.[10] They discussed the following two fundamental models for a self-manufacturing process:

1. The *unit growth* or *factory* model consists of a series of sub-assemblies, which collectively are able to manufacture and tassemble all sub-assemblies within the model until the surrounding resources are exhausted. As the name suggests, and as observed by von Neumann, any machine shop with sufficient equipment may be considered a self-manufacturing unit growth system.
2. Unlike the unit growth model, the *unit replication** model specifies that one device must be able to manufacture all of its own parts. Perhaps the most interesting facet of this model is that it potentially has the ability to be substantially more compact than the unit growth model, to the degree that such a kinematic machine could exist

*We shall call this the *unit reproduction* model from here on, in accordance with our defined terminology.

in every home. As yet, the autotrophic unit reproduction model has not been realised. One possible reason for this is that traditional manufacturing methods require tools to have one specific function (such as a lathe for cutting about an axis) severely limiting the potential designs that can be manufactured with a single kinematic machine. Therefore, the goal of achieving a self-manufacturing process based on the unit reproduction model cannot be realised until an extremely versatile manufacturing technology is realised.

2. The Genesis of RepRap

Sometimes the progress and the reporting of a project can obscure the train of thought that instigated the project. Typically, that train of thought was incomplete, or sometimes downright wrong. In this section, we attempt to set down as a matter of record the ideas that led one of us (Bowyer) to invent RepRap.

Understandably, the design of most practical artificial reproducers starts with proposed solutions to many technical problems of getting a kinematic machine copy itself. However, RepRap was not instigated in that way at all. RepRap was instigated by biomimetically considering extant naturally evolved *strategies* for reproduction.

Biologists categorise most bacteria, archaea, protists and plants as autotrophic because they are capable of self-nourishment using inorganic materials as a source of nutrients and using photosynthesis or chemosynthesis as a source of energy. However, almost without exception, all the natural species of reproducers in the world (including those in the previous sentence) depend upon other species in some way for their survival and successful breeding – by this light they are all assisted self-reproducers. A few lithophile micro-organisms can survive alone in what are essentially mineral environments, but their numbers are vanishingly insignificant compared with those of the interdependent species. Clearly, primordial organisms must have been completely autotrophic, but now that way of life has all but disappeared because the environment in which modern organisms have evolved consists, to a first approximation, entirely of other reproducers.

Yet research into artificial reproduction often concentrates upon making the reproducer as autotrophic as possible (like the lithophiles), and researchers regard this as an important aim. Clearly this aim *is* important for an extraterrestrial reproducer, but why so for a terrestrial one? Why try to follow a strategy that biology has all but abandoned? An

artificial reproducer designed to be interdependent with the natural reproducers that will make up its environment would be more likely to be successful.

Dependencies between species take one of the following three forms: *predation*, *commensalism*, or *mutualism*. Predation is well understood: lions eat antelope; antelope eat grass. Commensalism usually implies some sort of scavenging – lions and antelope are uninterested (though not ultimately disinterested) in what the grass does with their dung, their urine and their exhaled CO_2. Mutualism* implies a symmetry of dependencies giving benefit to both partners: the pistol shrimp digs a burrow in which it then lives with a goby fish; the shrimp is nearly blind and the fish warns it of danger.

This variety of dependencies prompts a choice in the design of an artificial reproducer: Which type of dependencies should our artificial reproducer exploit, and with which natural species? Beneficial options to people might include predation upon pests, commensal gathering of waste, or mutualism with a species whose welfare we wished to promote (such as an endangered or an agricultural one).

However, clearly the most interesting natural species with which our proposed artificial reproducer might exhibit a dependency is ourselves. This makes the choice more sharply cut: it would be foolhardy to make ourselves the prey for our artificial reproducer; having it collect our waste commensally might be useful, but the option most pleasing to our evolved senses of morality and symmetry would be to make ourselves a reproducing mutualist. In other words, we should make an artificial reproducer that would benefit from us, and we from it.

The most famous mutualism in nature, and the one that we all learn about first at school, is a reward for a service. Butler[11] said of this in *Erewhon***:

> Does any one say that the red clover has no reproductive system because the humble bee (and the humble bee only) must aid and abet it before it can reproduce? No one. The humble bee is a part of the reproductive system of the clover.

Moreover, he might have added, the bee is rewarded with nectar.

Mutualism between the flowers and the insects evolved about 140 mya in the late Jurassic period and is one of the most enduring phenomena in biology. For both sets of species it is an evolutionarily stable strategy corresponding to a particularly unshakable Nash equilibrium.

What service could our mutualist reproducer ask from us? Moreover, with what could it reward us?

It would seem sensible to play to the differing strengths of artificial kinematic machines and people. Artificial kinematic machines can make objects accurately, repeatably and tirelessly. In contrast, they fumble at manipulative tasks that would not tax a small child. People are exquisitely dexterous. (Aristotle called the human hand, *the instrument of instruments*.) But – though with practice people may

carve and mould beautifully – they cannot do so accurately, repeatably and tirelessly.

So our self-reproducing kinematic machine could be designed to manufacture a kit of parts for a copy of itself, and to need the assistance of people to assemble that copy (that is, it would be an *assisted reproducer* along the lines of NASA's unit-reproducer[10]). The people would be the humble bee, and the kinematic machine the clover. And what about the nectar? If the kinematic machine were sufficiently versatile to make its own parts, then chances are that it would also be able to make many other items useful to people. When it was not reproducing itself, it would be rewarding its assistants with a supply of consumer goods. This idea of a self-reproducing machine also making useful things for people is not new. It goes back through von Neumann to Butler. But we contend that regarding this as a form of biological mutualism and deliberately seeking to achieve that in order to position both reproducers at an evolutionary Nash equilibrium for each is a novel idea.

This was the genesis of the RepRap machine. It was designed to make its own parts to be assembled by people into another RepRap. The people would be driven to do this by the fact that the machine, when not reproducing, could make them all manner of useful products. It seemed (and still seems) likely that this would lead to a mutualist relationship between people and the machine that would inherit some of the longevity and the robustness of the evolutionarily stable strategies of the insects and the flowering plants.

Finally in this section, we note that flowers do not attempt some biological equivalent of copyrighting or patenting the "intellectual property" of their genomes. Such a genome builds the flower with the sole intent† of spreading itself with the most promiscuous fecundity possible. Any genome mutation that arose that – for example – attempted to extract some payment (like the nectar) in return for a copy of itself would clearly have a lower reproductive fitness. The nectar and the information are not in any way equivalent. The nectar is a real material resource. In contrast, the immaterial genome *information* has been arranged purely because of its success in copying itself as freely as it can, and any impediment placed in the way of that would be to its detriment. For this reason, it was decided to follow the principles of the free software movement and to distribute every piece of information required to build RepRap under a *software libre* licence that requires no royalty payments whatsoever. This would allow private individuals to own the machine, and to use it freely to make copies for their friends.

The RepRap machine is intended to evolve by artificial rather than natural selection; that is, to evolve as the Labrador has evolved from the wolf, rather than as the wolf has evolved from its ancestors. It is hoped that this evolution will come about by RepRap users posting design improvements on-line that may be adopted in future designs of the machine and then in turn downloaded by old and new users. That is why the General Public Licence was chosen as the RepRap licence, as that obliges people who improve the machine to make public their improvements under a similar free licence.[12]

---

*Sometimes called *symbiosis*, but this term (which can mean any species–pair relationship, including that between predator and prey) is being replaced by the more precise *mutualism* for mutually beneficial relationships.
**Famously, this novel is also one of the first places in history wherein appears the idea of an artificial reproducing machine.

†Of course, no genome has *intent*. But they all behave as if they do.

3. The First RepRap Machine

In order for the RepRap project to progress, the rather abstract reasoning laid out in the sections above had to give way to some down-to-earth engineering. This occupies the rest of this paper.

The first engineering decision was to use rapid prototyping* as the manufacturing technology for RepRap as opposed to – say – the CNC milling. The reasons for this were threefold:

- Rapid prototyping requires very low forces to create solids, unlike machining.
- Of all the manufacturing technologies, rapid prototyping is the easiest to control automatically by computer.
- It is the closest current technology to the "maximum versatility" specified by the 1980s' NASA report.[10]

The massive cast bases required by machining centres to overcome cutting forces and vibration would be a significant impediment to self-reproduction; this was the basis of the first reason. It is true that cold-casting of – for example – concrete might have reduced this impediment to a certain extent, but that would still leave the fact that cutting (especially of complicated re-entrant shapes) requires very sophisticated toolpath-planning algorithms. In contrast, all rapid prototyping requires algorithmically is the ability to compute a sequence of planar slices through a geometric model of the part to be manufactured and to fill each with a hatch or similar pattern (see the second reason above). These are straightforward.

At around this stage it was decided (as mentioned in the Introduction) that any parts that the RepRap machine could not make for itself had to be cheaply and widely available to maximise the ease – and hence probability – of reproduction.

Having chosen rapid prototyping, the next decision that needed to be taken was which of the extant processes to use, or whether it would be necessary to invent a new one.

This decision was made by elimination: any process needing a laser was rejected owing to the unlikeliness of being able to use rapid prototyping to make a laser, and the fact that they are expensive items. This removed selective laser sintering and stereolithography from consideration (and also electron beam melting for very similar reasons, though that does not actually need a laser).

Similarly, any process needing inkjet print heads was rejected. Again, it was thought unlikely that the machine would be able to make these for itself, at least initially. Though ink-jet heads are relatively cheap, they are constrained by the fact that their manufacturers have a business model of discounting the printers that use the heads and then putting a big mark-up on the heads themselves. Also, they sometimes put chips in the heads to prevent their being re-filled, and adopt other restrictive strategies that make inkjet an unattractive technology for a machine intended to reproduce as freely as possible. These facts removed ink-jet printing from consideration.

The project was then left with laminated object manufacturing and fused-filament fabrication from the extant technologies. Laminated object manufacturing was attractive owing to its simplicity and the ubiquity of its working material: paper. But fused-filament fabrication[13] offered the possibility of being able to build with multiple different materials. This in turn offered the significant potential advantage of being able to have the machine make a larger proportion of its own components than could be created out of just one material. This, combined with the fact that it was conjectured that fused-filament fabrication could be implemented using low-cost garden-shed methods,** led that to be chosen for RepRap. Thus, it was not necessary to devise a new rapid prototyping technology.

Figure 3 shows the first production the RepRap machine. There were experimental machines made before this to try out various ideas, but this is the first model used to make copies of its own rapid prototyped parts. It is a stepping-motor-driven Cartesian robot consisting of an open frame made from M8 threaded steel rods held together by rapid prototyped parts and M5 screws. The build base upon which parts are made is cut from 12-mm medium-density fibreboard (MDF). Virtually all the parts that the machine cannot make for itself other than the electronics and the motors can be obtained from an ordinary high-street hardware shop. The horizontal x- and y-axes (which need to move comparatively fast – typically with a feed rate of 3000 mm/min) are driven directly by toothed timing belts. The z-axis only moves by a small distance when one layer of production is finished and the next is about to start; aside from this and the need to return to its home position at the start of a new build, the z-axis is quiescent. It was thus decided to use a screw-drive for that. The MDF build plate has eight M8 nuts attached at its corners that are driven up and down by four M8 threaded bars synchronised by another timing belt. The nuts are in pairs held apart by springs to eliminate backlash.

The entire machine is designed to work from a single 12-volt power supply as this can be cheaply obtained by using the power supply from an old PC. It also means that the machine would work off a car battery, where no main electricity was available. The running machine consumes about 60 W of power.

Given the choice of the fused-filament fabrication, a key – indeed *the* key – part of the machine is its polymer-extruder head. This is described in the next section.

3.1. Extruder design

Before a fused-filament fabrication extruder was designed, a polymer had to be chosen for it to extrude. It was decided to use polycaprolactone (PCL) initially because:

- It has a very low melting point (about 60 °C).
- It is strong (comparable to nylon).

As will be seen below, this polymer was subsequently abandoned for several reasons. But at the start, the low melting point (implying ease of heating) was thought to be a critical factor.

---

*The term "additive fabrication", which is now becoming a carryall for the totality of rapid prototyping and 3D printing technologies, was less current at the start of the RepRap project (2004). Then "rapid prototyping" was by far the more popular term. Its subsequent decline has tracked the change in use of these technologies from prototyping towards production.

**A conjecture that the project has subsequently proved correct.

Fig. 3. RepRap version I "Darwin". This is the first production RepRap machine. Its rapid prototyped parts (white, blue and green) were made in a Stratasys Dimension commercial RP machine. The cube of the machine has side lengths of about 500 mm. This machine was built in May 2007.

A fused-filament fabrication extruder has the following two main sub-assemblies:

1. The transport; this forces a filament of the polymer into. . .
2. . . .the melt chamber and nozzle.

It was decided to standardise on a 3-mm diameter* filament as a feedstock because this dimension is commonly chosen for plastic welding rod, which is very widely available.

3.1.1. Polymer transport. Figure 4 shows the first polymer transport mechanism designed for RepRap. A 12-volt DC geared electric motor (A) drives a stack of pinch wheels (C) through gears (behind the device). The pinch wheels are actually the heads of M4 cap screws; the outside knurling on these gave grip. The threaded rods at the front allow the pinch-wheels to move together, thereby increasing the pinch force. The number of pinch wheels in the stack could be varied. The 3-mm polymer filament (B) was driven down into a melt chamber and nozzle at D (see figure). The control electronics are at E. All the white parts in the photograph except D were rapid-prototyped.

This device worked, but it was heavy and complicated. The large number of pinch wheels in the stack (4) was needed because of the very low friction coefficient between the pinch wheels and the PCL (despite the knurling).

Shortly after this device was made, one of us (Olliver) came up with a much simpler design. In it, a threaded rod is forced against the 3-mm polymer filament, which runs in a channel. As the thread turns, it forces the polymer downwards. This design, which has only one moving part, gives exceptional


Fig. 4. The first RepRap polymer transport mechanism. It is shown diagrammatically on the right: four pairs of pinch wheels are gear-driven by a motor and force a 3-mm polymer filament (blue, then orange, then red to indicate the rising temperature) down through a PTFE insulator then through a heated brass barrel with a nozzle. The molten polymer emerges at the nozzle's tip.

grip even against the slipperiest polymers, and has a very high mechanical advantage, giving a large potential extruding force.

---

*1/8″ diameter in the United States. 1/8″ − 3 mm = 0.175 mm, a difference easily accommodated by all the extruder designs that RepRap has used.

Fig. 5. The last screw-driven RepRap polymer transport. All green parts are rapid-prototyped.

The design of the screw-driven extruder went through several iterations, finally ending up as shown in Fig. 5.

The geared drive (at the top) is offset, with its torque being transmitted by a flexible coupling (a short length of steel hawser) – the grey curve just to the right of the white polymer filament. This arrangement allows a straight run of filament down the device, which it was thought might be useful for brittle or stiff materials. However, this straight run was never, in fact, used (the polymer was always flexible enough to feed in at an angle). Further, the flexible drive was a weak point in the design, as it tended to fatigue after being used for about 50 h. The springs at the back of the device set the force between the screw thread and the polymer, and allowed some compliance as slight changes in polymer diameter moved through the device.

As these developments were taking place, the temperatures that could be easily achieved in the melt chamber (see below) were rising because of design improvements. This meant that it was possible to abandon PCL, which didn't just give problems because of its low friction, but was also very sticky and stringy as it was being extruded. These shortcomings led to low-quality built parts. Acrylonitrile butadiene styrene (ABS) was adopted instead. This gave much better build quality (as it was more paste-like upon extrusion and less viscous). ABS also allowed a return to a much simpler pinch-wheel transport mechanism, as, being harder and exhibiting higher friction, it was easier to have a firm grip.

Figure 6 shows the extruder design used in RepRap Version II "Mendel". A NEMA 17 stepping motor with a knurled shaft pinches the filament against a ball-race. The motor has a 5-mm diameter shaft. With this, a motor torque rating of 0.13 Nm is quite adequate to drive the filament with enough force for reliable extrusion. The device has a single rapid-prototyped part (silvery white). The stepping motor allows exquisitely precise metering and control of the extruder flow. The four screws holding the motor are in slots. This allows the gap between the motor's shaft and the ball-race to be adjusted easily. For 3-mm hard-polymer filaments, a 2.5-mm gap works well. This is easily set by putting the shank of a 2.5-mm drill-bit in the device, sliding the motor so that the bit is just trapped between the motor's shaft and the ball-race, tightening the four screws, then withdrawing the drill-bit.

3.1.2. Melt chamber and extrusion nozzle. The requirements for the melt chamber and nozzle were that it should:

1. Be cheap and easy to make;
2. Be compact;
3. Work reliably after repeated heat-up and cool-down cycles; and
4. Not conduct excessive heat to the rest of the machine.

Point 4. was particularly important, as the rest of the machine would be made from the polymer that the melt chamber would be melting.

Figure 7 shows the first design of melt chamber and nozzle. The 3-mm filament enters on the left, and is extruded from a 0.5-mm diameter nozzle on the right.


Fig. 6. Left: The newer pinch-wheel extruder transport mechanism (with melt chamber and nozzle pointing to the right). Right: The motor's shaft (which is knurled-the small circle) pinches the filament against a ball-race (large circle) and drives it down through a PTFE insulator into a heated chamber and out through a nozzle.

Fig. 7. The first melt-chamber and nozzle design. The small squares have 1 mm sides.

The white cylinder on the left is a 16-mm diameter polytetrafluoroethylene (PTFE) tube. The internal diameter of the hole running down it is 3.5 mm, which was found to work well with the 3-mm filament (more on this below). The right-hand end of the PTFE has an M6-threaded hole extending to a depth of 15 mm, into which a length of drilled brass M6 studding has been screwed tightly. Again, the drilled hole in the studding is 3.5 mm in diameter.

The nozzle at the right end is turned brass. A nozzle of 0.5-mm diameter was chosen, as that is the smallest hole that can be drilled easily and is also a good compromise between the machine's being able to make fine details (see below) and its not taking too long to fill a large volume.

The heating element (the left-hand pair of wires) was a length of fibreglass-insulated nichrome wire with a resistance of around 10 Ω. This was wound in the grooves of the M6 thread, giving good thermal contact. Temperature sensing was done with a 10-KΩ glass-bead thermistor (right-hand wires). Both these were held on using JB Weld commercial high-temperature epoxy, which is rated up to 315 °C.

This design worked well, particularly because PTFE has a very low thermal conductivity, which kept the rest of the machine cool. But it suffered several shortcomings:

1. The PTFE was held in the polymer transport mechanism by a screw clamp. It tended to slip free of this because PTFE has a low friction coefficient (in contrast, a good thing from the perspective of the polymer filament being forced down the middle of it).
2. A resistance of 10 Ω was too high, giving too low a heating power at 12 volts.
3. The JB Weld tended to become friable and was easily damaged after being subjected to a large number of heating and cooling cycles.
4. The 10-KΩ thermistor was not very accurate above about 200 °C.
5. PTFE is a rather soft plastic. This meant that sometimes the inner tube swelled under the pressure of the filament being forced into the heated brass tube, leading to an aneurysm of extrudate. The device continued to work while heat was applied, but the swelling could cause blockage when the device had cooled and was restarted.

The problem of the PTFE slipping in its clamp was solved by turning a series of grooves on the left-hand end of it. It was then epoxied into the transport mechanism with the epoxy keying into the grooves. This held completely firm.

The heater resistance was reduced to 6 Ω in subsequent versions. This gave an ample heating power of 24 W.

The JB Weld problem was solved by replacing it with fire cement (which is used to seal the flues of central-heating boilers). This is rated up to 1250 °C. This fire cement was subsequently replaced with Kapton-brand high-temperature sticky tape; this made extruder assembly simple, clean and neat.

The thermistor was replaced by a 100-KΩ one. This corrected the loss of accuracy at high temperatures.

A number of changes were tried to reduce the problem of PTFE swelling, and this has now been completely eliminated (see below). Improvements tried included using polyaryletheretherketone (PEEK) instead of PTFE. This is mechanically much stronger, but the internal friction is higher, and so jams can occur. It is also not a good thermal insulator, and so it has to be made longer to achieve the same cooling effect. We also tried using a PTFE sleeve inside a PEEK outer jacket.

One of us (Palmer) has had considerable success by abandoning high-temperature polymers as thermal barriers altogether and instead using stainless steel plus a heat-sink (Fig. 8). Here, the heating is provided by a 6.8 Ω resistor embedded in an aluminium block heating a copper nozzle. This is connected to the polymer filament transport by a stainless steel tube cooled by a heat-sink and small fan. As stainless steel has a low conductivity (for a metal), it is possible to maintain a high-temperature differential between the two ends. But the device can still jam. This problem can be solved by adding a very slight taper into the hole down the centre of the device from narrow at the top to broader at the bottom. This means that as soon as the cold polymer begins to move when the device is started, it detaches from the walls and is subjected to almost no friction. This is a very reliable design.

The current standard RepRap extrude nozzle is also very reliable. In it, the brass nozzle is made larger in diameter than before, and encases the end of the PTFE. When pressure

Fig. 8. The stainless steel heat-sink melt-chamber and nozzle.

Fig. 9. The latest standard RepRap nozzle design partly disassembled. The PTFE ends in a narrow section that screws inside the brass. The PEEK bracket is the part with three holes. The brown material at the top of the PTFE is Kapton tape to ensure a snug fit in the extruder transport mechanism.

causes the PTFE to swell, this forces it against the brass, constraining it and creating a better seal. A small PEEK bracket attaches the brass nozzle to the rest of the extruder mechanism so that none of the extrude force is born by the PTFE (see Fig. 9).

4. The RepRap Build Materials
Polycaprolactone was quickly abandoned as a building material after initial experiments for the reasons outlined above. Though it is extremely tough, it is also quite flexible. For building the RepRap machine, something more rigid – that is with a higher Young's modulus – was desirable. The Young's modulus for PCL is about 1 Gpa, and that for ABS is about 3 Gpa.* ABS was chosen next, as it is inexpensive and very widely available.

Figure 10 shows a typical RepRap build of a reasonably large (80 mm × 60 mm × 30 mm) component using ABS extruded at 240 °C. As can be seen, the contraction of the ABS on solidification has caused the bottom of the part to curl away from the base upon which it was built. For smaller parts, this is much less of a problem, but here it is significant. Commercial fused-filament fabrication machines solve this problem by running the entire process in an oven at a temperature of about 70 °C. This eliminates curl-up with ABS completely. However, since the point of RepRap is that it should be easy for any technically competent person to construct and to run the machine at home in order that its dissemination should be as wide as possible, running in an oven was not really an option. We did experiment with enclosing the build in a roasting bag (intended for joints of meat) and filling it with hot air from a hair dryer. This worked very well and completely eliminated the problem. But the bag was tricky to set up, it tended to get tangled in the moving mechanism and a lot of work was required to unload the built parts from within it.

Parallel to the emergence of this curling-up problem, one of us (Olliver) was experimenting with the use of polylactic acid (PLA) as a RepRap building material. PLA is slightly harder and more brittle than ABS. These experiments were not prompted by a need to solve distortion in builds, but by ecological and socio-political considerations. PLA is made from plants and is biodegradable. Its use in RepRap would thus be more environmentally benign than the use of an oil-based polymer, such as ABS. Indeed, as PLA is made from plant matter, building durable objects from it that are kept and *not* bio-degraded would lock up atmospheric CO_2; this would be environmentally positive as opposed to merely neutral. In addition, it is not too hard for people having access to a small starch crop to make their own PLA.** This means that such a person with a RepRap machine would not only be able to make the rapid-prototyped parts to reproduce more RepRaps and to make other goods, but would also be able to do so with a home-grown polymer supply that was also self-reproducing. And they would be reducing greenhouse gas as they worked.

Serendipitously, it transpires that PLA suffers minimally from contraction problems on cooling, even when builds are conducted in a room-temperature environment.

---

*Of course, these values are very dependent on the degree of polymerisation.
**There is one step in the synthesis that is not straightforward: the lactic acid that results from the fermentation of starch has to be dried to better than one-part in 10 million by weight of water. We have succeeded in doing this by simply passing dry nitrogen over it for half an hour with heating. We conjecture that it would also be possible to do it with air that had been dried by passing it through calcium chloride and then heated. The calcium chloride could in turn be reused after drying by heating.

Fig. 10. A RepRap extruder component made in a RepRap machine from ABS. Note the curling at the base.


Fig. 11. A clock made by a RepRap machine from PLA. The face was made from 12 separate segments, each a longest dimension of 140 mm. (The rather scruffy octagonal backplate is the MDF-built platform from an old experimental prototype made to try out ideas for the machine in Fig. 3.)

Figure 11 shows a clock made on a RepRap machine from PLA extruded at 190 °C. The twelve hour-segments were made separately. Each is 140-mm long in its radial direction. With other polymers significant curling would be expected when building such an object, but PLA gives almost completly flat parts.

In addition to its positive environmental characteristics and its lack of distortion on building, PLA welds to itself very strongly when one fused-filament fabrication layer is being deposited on top of another. Parts made from it are less susceptible to delamination under stress in the vertical-build direction than those made from ABS. It is an almost perfect material for the fused-filament fabrication process. Its disadvantage (aside from the fact that it is only commercially

available from a small number of companies) is that it tends to "string" slightly, leaving filaments sticking out from builds where the extrude head has left the build for an in-air movement. This problem can be eliminated by reversing the extruder by a short distance at the end of each completed deposition before an in-air movement. Any remaining stringing can be removed manually with a blade in a few seconds after the build is completed.

5. RepRap Software

A number of people have created software to drive RepRap machines. Here we describe our program written in Java to take STL files of the objects that RepRap is to make, to slice those, to compute infill patterns and to save the results as G-Code NC control files (the format that RepRap expects). For convenience, the same program can also take such pre-computed G-Code files (from any source) and queue them to a RepRap machine.

Figure 12 shows this software being run. On the left is a window that allows the user to control the attached RepRap machine interactively for setting up, testing and experimentation. On the right is a view of the machine's build base with a set of components (in fact, parts of the RepRap machine itself) loaded for manufacture.

The user can interactively view these components from any point, move around parts with a mouse, insert new parts and delete unwanted ones. The whole collection can be saved in a single file for future use, if required.

In order to generate the control codes for the RepRap machine, the software takes a series of horizontal slices through the STL files making up the parts to be built. As STL files consist entirely of triangles, this is simply a question of working out the line segments that result from the intersection of those triangles with the slicing plane. These are then joined into polygons by finding nearest end-points of lines.

The software then uses Tang and Woo's Algorithm[16] to convert those polygons into an interim CSG representation as intersections and unions of linear half-planes of the form

Fig. 12. The user's view of the Java RepRap software.

$Ax + By + C <= 0$. This is done because a CSG representation is much more robust than a boundary representation. If the STL input is wrong, the shapes will be wrong (that is unavoidable). But with CSG they are always topologically consistent.

The CSG polygons are then converted to bitmaps at the finest resolution of the RepRap machine. This is done because Boolean operations will have to be done on the slices represented, and this is simplest if they are in bitmap form. It might be thought that the CSG representation would be ideally suited to Boolean manipulation, and indeed our software worked that way initially. But problems arose with repeated Boolean operations on patterns that were very similar but (because of rounding and other factors) not identical. Bitmaps completely eliminate these problems, and – an added advantage – can be processed very fast.

The slices are computed from the top to the bottom of the parts to be built. This allows the easy calculation of support material: layer L_{n+1} needs the following support, S_n, from layer L_n:

$$S_n = L_{n+1} - L_n.$$

Moreover, the layer pattern at n potentially requiring support at layer $n - 1$ would be

$$L_n \leftarrow L_n \cup L_{n+1}.$$

As can be seen, this calculation is most easily facilitated from top downwards.

In fact, the software maintains a cache of layer patterns in a ring buffer because it is useful to know the patterns a few above and a few below the current slices for the purpose of computing fine infill at the surface when coarse infill is being used for interiors: a horizontally exposed surface facing either up or down has to be fine-filled to get a well-finished part. The cache ensures that slices are computed only once, while not making the excessive memory demands that keeping a complete slice record would entail.

Once a slice has been computed, the outline polygons are found using a pixel edge-finder plus a filter to collapse straight sections of many pixels back to a single line segment (the filter picks out runs of pixels that would be generated by a Bressenham DDA between their endpoints, and replaces them with those endpoints).

The infill hatch is again generated using a Bressenham DDA. When the DDA goes from an empty pixel to a filled one that is the start of a hatch segment; the reverse is its end. The resulting hatch lines are joined by traversing sections of polygon boundary between their ends to produce a set of zig-zag lines filling the polygon.

The Java RepRap software can deal with multiple objects, each made from multiple materials. It contains all the internal controls to allow these to be outlined, filled and supported in a completely general way. Thus, an object made from material A can be outlined for one or more times going inwards (to make a thin or thick boundary); it can then be filled with material A with any hatching width, with the hatch angles changing between layers; it can be supported by material A laid down thinly to make it easy to break off after completion. Alternatively, the above can be done with differing materials A, B and C. Finally, multi-material objects can be specified, and their peripheries and interiors can be made in the same completely controlled way.

6. Reproduction of RepRap Machines

The first reproduced RepRap machine was made in 2008. The parent machine is the one shown in Fig. 3. All the rapid prototyped parts of the child machine were made in PLA by the parent machine, except for one grandchild part (a timing-belt tensioner) that the child machine made for itself. That grandchild part was the first part made by the child. It took about 20 minutes to make, and was finished at 1400 h UTC on 29 May 2008 at Bath University in the United Kingdom.

The child machine was within tolerance of the parent, and worked just as well. RepRap is thus a *kinematic-assisted self-replicating, self-manufacturing machine*. The design of the machine includes screw and other adjusters to allow a child

machine to be set up to produce parts as accurately as its parent (in just the same way as conventional machine tools are adjusted). RepRap assisted replication is thus not subject to degeneracy. However, its accuracy does, of course, depend upon the measurements made to achieve the adjustments just mentioned, and the measuring instruments (typically digital callipers) are – like the motors and so on – external to the machine. It would doubtless be possible to set up a calibration scheme depending on light interference or similar that could be a part of RepRap itself. But this would be a heroic effort to overcome a simple externality and, as we shall show in the next paragraph, RepRap has many of those.

The fraction of the parts by count of the RepRap machine that it makes for itself is 48% ignoring fasteners (that is, nuts, bolts and washers – if these are included, then they are 73% by count of the machine and the RP parts drop to 13%).[14] It would be straightforward to make a much simpler and rather cheaper machine of an almost identical design by gluing the parts together rather than using nuts and bolts. The machine could easily make small cylindrical locating lugs to fit in the bolt holes and to hold the gluing parts in the correct relative positions while the glue set, but that would make it harder to replace parts in the machine and to service it.

Many people have downloaded the RepRap designs from the project website and used them to make RepRaps in commercial rapid prototyping machines. In one case (father and son constructors Bruce and Nick Wattendorf) the design was cut from wood, assembled and made to work successfully.

The RepRap community calls free open-source rapid prototyping machines that cannot make a significant fraction of their own parts, but can make parts for RepRap machines, *RepStraps* (from *bootstrap*). Many private individuals (including one of us – Palmer) have made RepStraps in order to make RepRaps subsequently. Several companies have been formed to make and sell RepStraps (for example, Bits From Bytes in the United Kingdom, and MakerBot Industries LLC in the United States). As required by the GPL, these companies' RepRap-derived designs are being distributed free and open-source.

Developed independently of RepRap, though inspired by it, the Fab@Home fused-filament fabrication machine from Cornell[15] is laser-cut, as are the Bits From Bytes and MakerBot machines. Fab@Home works primarily with a variety of pastes (including UV-setting resin), though a thermoplastic extrusion head has also been attached to it. The Bits From Bytes machine is almost identical in size and capability to RepRap I "Darwin". The MakerBot machine is smaller than "Darwin", but is about the same size as RepRap II "Mendel" (see below), though with a smaller build area. Both these commercial machines tend to be used with ABS, though some of their users have also built with PLA. Both Bits From Bytes and MakerBot can make all the parts for RepRap. We have heard no reports of this having been done with a Fab@Home, though there seems to be no technical impediment to its happening.

Increasingly, people are using their RepRap machines to make sets of RepRap parts for others, as the project plan intended.

The first complete set of rapid prototyped parts made in one RepRap machine that were sold to other would-be RepRap constructors were made by Wade Bortz from Canada. The child machine's parts were bought by Liav Koren and Michael Bartosik of Toronto, who paid one case of Upper Canada Dark Ale for them*. Given the analogy with the payment of nectar for flower reproduction discussed in Section 2, the case of beer is particularly appropriate.

Owing to the free distribution of the machine it is difficult to make a worldwide estimate of the number of RepRaps and RepStraps there are, but the sale of electronic kits for the machine (which are also produced commercially) sets a lower limit of 3000 machines. However, some people construct their own electronic kits rather than buying from market. About 4500 machines would seem to be a conservative estimate of the total population at the time of writing this paper (i.e., in 2010).

The RepRap website invites builders of the machine to mark their location on a map. Figure 13 shows this map as in 2010. Only a small fraction of builders have placed themselves on it, but it gives an interesting (if self-selected) sample of the distribution of the machines.

7. Changes Made to Produce RepRap Version II

The authors and their many RepRap colleagues around the world have now finished the design and commissioning of the latest RepRap machine: RepRap Version II "Mendel"** (Fig. 14). This incorporates many lessons that were learnt from Version I; in particular, lots of improvements and suggestions from the worldwide RepRap community that were posted on-line in the project's forums have been included in the design.

Table I gives a comparison between RepRap Version I "Darwin" and RepRap Version II "Mendel". The cost-of-materials-to-build figures are for a single purchase of all that is required to build one machine from end retailers – it takes no account of bulk discounts and wholesale transactions.

The designs of both machines allow their sizes and working volumes to be changed simply by cutting longer or shorter rods to make up the framework, so the values for both of these are nominal.

Darwin may be carried a short distance by one person with some difficulty. "Mendel", in contrast, can be swung in one hand like a bulky briefcase, and is easy to carry anywhere. This has turned out to be a surprisingly important improvement – portability makes it more convenient to work with the machine in many contexts.

The percentage of the machine that "Mendel" makes for itself has remained constant in comparison to "Darwin", despite a significant number of rolling element bearings being incorporated into the design to give robustness. Furthermore, some users in the community have replaced these bearings with plain bearings. In this case, the self-manufactured percentage rises to 57%. It is anticipated that the number of self-manufactured parts will rise further once the multiple write-heads are finished (see below).

*At the time of writing, the lowest cost of non-open source rapid prototyping machine (the V-Flash made by 3D Systems) was about $ 9900.
**It has been decided to name RepRap versions after distinguished biologists.

Fig. 13. Some of the world's RepRap users in 2010 (courtesy: Google Maps). (A few users who do not wish to reveal their location have put themselves in the ocean.)

The deposition rate of 15 mL/h is the volume extruded by the extruder. In common with poprietary fused-filament fabrication machines, RepRap does not usually build parts completely solid – there are some air inclusions. With RepRap the degree to which this happens is completely under the control of the user. It is possible to build parts very fast with a sparse honeycomb interior, or more slowly with a dense interior. Unlike the commercial fused-filament fabrication machines (which leave micro-voids on their densest settings), RepRap also allows interiors to be built fully dense. This slightly reduces the build quality, but allows the making of gas- and water-tight parts. With the nominal settings of the machine, the 15 mL/h deposition rate becomes 19 mL/h of built object.

We have not gathered reliability statistics on the machine as yet (it is very easy to make spare parts, so repair is not a problem), but a "Mendel" machine will typically run for several hundred hours without going wrong, though it might need occasional fine adjustments in that time (for example to the bed-height zero-position, or the filament pinch drive mechanism).

As has been mentioned above, an 0.5-mm nozzle diameter was chosen as a compromise between ease of manufacture, speed of deposition and fineness of feature resolution. However, it is perfectly possible to make nozzles of different diameters without changing any other aspect of the RepRap machine's design. Drilling very small holes is difficult, of course. But as Jens Kaufmann of Heriot Watt University has suggested that it might be possible to make fine nozzles by running copper sulphate solution through an 0.5-mm brass nozzle and electroplating copper onto its inner surface – an experiment that we have yet to try. With the 0.5-mm nozzle, the smallest features that we have produced (gear teeth for the machine itself) are of about 1.5-mm in size.


Fig. 14. RepRap Version II "Mendel". This version is smaller, lighter and simpler than Version I, but it has a larger build volume.

Table I. Comparison of "Darwin" and "Mendel".

	RepRap version I "Darwin"	RepRap version II "Mendel"
Cost of materials to build (€)	500	350
Percentage self-manufactured*	48	48
Size (mm^3)	600 (W) × 520 (D) × 650 (H)	500 (W) × 400 (D) × 360 (H)
Weight (kg)	14	7
Build envelope (mm^3)	200 (W) × 150 (D) × 100 (H)	200 (W) × 200 (D) × 140 (H)
Deposition rate (mL/h)	15	15
Positioning accuracy (mm)	0.1	0.1
Nozzle diameter (mm)	0.5	0.5
Volume of RP parts to build (mL)	1200	1110
Power supply (W)	12 V × 8 A	12 V × 5 A
Interface	USB/G-Codes	USB/G-Codes
Host computer	Linux, Windows or Mac	Linux, Windows or Mac

We are currently working to add the final part of the "Mendel" design: a multi-extruder head that would allow the machine to print with multiple materials. The reader will recall that this potential was one of the initial reasons for choosing fused-filament fabrication for RepRap. We also have a paste extruder in the later stages of development (driven by compressed air from a fizzy-drink bottle acting as a reservoir and charged using a car tyre pump). This should allow RepRap access to a wide range of materials already proved in fused-filament fabrication by the open-source Fab@Home project mentioned above.

8. Conclusions

There is no space here to go into every last reason for adopting certain designs for parts of RepRap and rejecting others, nor to describe all the many alternatives that were experimented upon and not adopted because of the results of those experiments. All this information is, however, available in copious detail on the RepRap blogs, forums and wiki (http://reprap.org).

At the start of 2008 four RepRap machines existed, all made on commercial rapid prototyping machines. Two and a half years later we conservatively estimate that there are about 4500 derived machines all over the world. We have no way of telling how many of those are replicated RepRap machines, and how many are non-replicating RepStraps for making further RepRaps. However, judging by the large number of requests for the fused-filament fabricated parts for the "Mendel" design when it was released, lots of people want to make, to use and to distribute their own assisted replicator.

One of the members of the RepRap project (Zach Smith) has set up a website where anyone can upload and download free designs for consumer goods to be printed by RepRap and other rapid prototyping machines (http://thingiverse.com). This is a considerable success, with many new designs being added daily.

RepRap works well. Nevertheless, even a poor reproducer that is out in the world freely parenting children must improve by Darwinian selection, and so should eventually overtake even the most exquisitely designed reproducer that stays in the laboratory.

*Ignoring nuts, bolts and washers.

The reader will note that this paper contains no details of the future direction of the RepRap project. The reason for this is that the authors are no longer in control of it. They will contribute to future developments, but increasingly those developments come from people in the RepRap community, and what they will do is well nigh impossible to predict.

All human engineering can be considered to be a vast unit-growth reproducer that copies itself with the assistance of – and with benefit to – humanity; a grand version of von Neumann's well-equipped machine shop. RepRap is moving towards compressing as much of that idea as possible into a unit-reproducer that one human may carry in one hand, and may freely copy for their friends.

References
1. C. E. Shannon, "A mathematical Theory of Communication," *The Bell Syst. Tech. J.* **27**, 379–423, 623–656 (Jul.–Oct. 1948). (http://tinyurl.com/f4two)
2. R. A. Freitas and R. C. Merkle, *Kinematic Self-Replicating Machines* (Landes Bioscience, Georgetown, Texas, 2004). (http://www.molecularassembler.com/KSRM.htm)
3. J. Von Neumann, *Theory of Self-Reproducing Automata* (University of Illinois Press, Illinois, 1996).
4. A. G. Cairns-Smith, *The Life Puzzle* (University of Toronto Press, Toronto, Ontario, 1971).
5. F. J. Dyson, "The world, the flesh, and the devil. Self-reproducing machinery," **In**: *The Third J. D. Bernal Lecture*, Birkbeck College, London (May 16, 1972).
6. L. S. Penrose, "Mechanics of self-reproduction," *Ann. Hum. Genetics* **23**, 59–72 (1958). (http://vx.netlux.org/lib/mlp01.html)
7. M. Moses, A Physical Prototype of a Self-Replicating Universal Constructor *Masters Thesis* (Mechanical

Engineering, New Mexico: University of New Mexico, 2002). (http://home.earthlink.net/ mmoses152/SelfRep. doc)

8. J. Suthakorn, A. B. Cushing and G. S. Chirikjian, "An Autonomous Self-Replicating Robotic System," *Proceedings of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM 2003)*, Kobe, Japan (December 2003). (http://tinyurl.com/ylfgr9b)

9. V. Zykov E. Mytilinaios, B. Adams and H. Lipson, "Self-reproducing machines," *Nature*, **435**, 163 (2005). (http://tinyurl.com/ygozg37)

10. R. A. Freitas and W. P. Gilbreath, "Advanced Automation for Space Missions," *Proceedings of the NASA Conference Publication CP-2255 (N83-15348)*, Summer 1980 (1982). (http://www.islandone.org/MMSG/aasm/ accessed 22 November 2010)

11. S. Butler, *Erewhon*, (1872). (http://www.gutenberg.org/etext/ 1906)

12. R. Stallman *et al. The GNU General Public License* (Free Software Foundation, Inc., Boston, MA, 1991). (http://www.gnu. org/licenses/gpl.html)

13. S. S. Crump, 1989. US Patent 5121329: *Apparatus and Method for Creating Three-Dimensional Objects* (Stratasys Inc., Eden Prairie, MN, 1989) (http://tinyurl.com/ybtrkft)

14. E. A. Sells, Towards a Self-Manufacturing Rapid Prototyping Machine *Ph.D. thesis* (Bath, UK: University of Bath, 2009). (http://tinyurl.com/ygjgva7)

15. E. Malone and H. Lipson, "The Factory in your Kitchen," *Proceedings of the MCPC 2007 World Conference on Mass Customization & Personalization*, MIT, Cambridge, MA, USA (October 7–10, 2007).

16. K. Tang and T. C. Woo, "Algorithmic aspects of alternating sum of volumes. Part 1: Data structure and difference operation," *Comp. Aided Des.* **23** (5): 357–366 (1991). (http://tinyurl.com/yg4wucm)