G R E S H A M

COLLEGE

THE PRACTICAL FRACTAL

A Lecture by

PROFESSOR IAN STEWART MA PhD FIMA CMath
Gresham Professor of Geometry

2 October 1996

Gresham Geometry Lecture
2 October 1996

The Practical Fractal

About twenty-five years ago Benoit Mandelbrot, then a scientist working for IBM, became aware that there was a common thread running through his work. He had been studying all sorts of different problems — the stockmarket, the amount of water in rivers, interference in electronic circuits. The common thread was that each problem had intricate structure on all scales of magnification. If you graph price movements in the stockmarket on a monthly basis, you get a rather irregular curve with lots of ups and downs. If instead you look on a weekly basis, a daily one, an hourly one, or even minute by minute, you *still* get a rather irregular curve with lots of ups and downs. The same goes for the water flowing in a river, or the changes in current in a noisy electronic circuit.

Mandelbrot decided that this kind of structure needed a name, and he invented one: **fractal.** A fractal is a geometric shape that has fine structure on all scales of magnification. Most of the familiar forms in the natural world are fractals. A tree, for example, has structure on many scales: trunk, bough, limb, branch, twig. So does a bush, a fern, or a cauliflower. A lump of rock looks like an entire mountain in miniature; a small cloud looks just as complicated as a big one if you view it close up; the surface of the moon is covered in craters of all sizes; the coastline of Britain has promontories and inlets of all sizes.

The traditional shapes of mathematics do not behave like this. If you magnify a sphere, then its surface becomes flatter and flatter, resembling a featureless plane.

In the ensuing decades, fractals flourished. They became widely recognised, thanks to the remarkable graphics that they create. Here's an example:

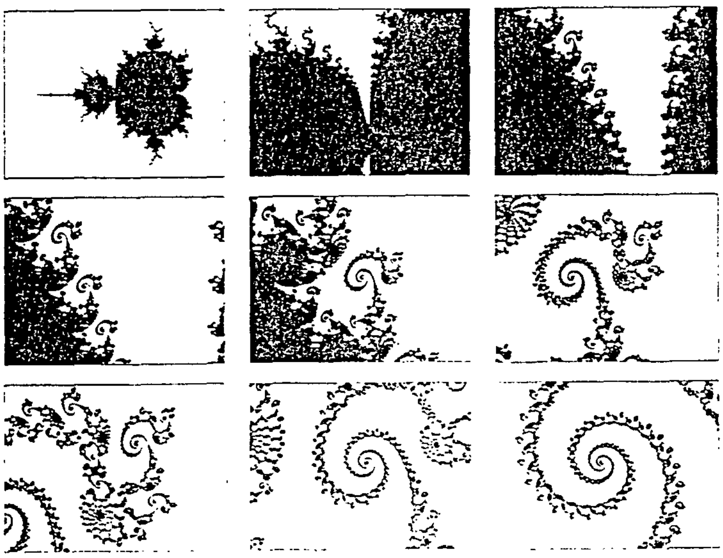

Fig.1 Part of the Mandelbrot set: the archetypal fractal.

In recent years several commentators have deplored the absence of serious applications for fractals, declaring them to be no more than pretty pictures.

This is absolutely untrue.

This lecture describes two of the practical uses to which fractals — and of course the associated mathematical machinery — have been put. The first is Michael Barnsley's technique of FRACTAL IMAGE COMPRESSION. The other is the FRACMAT machine for quality control of spring wire, invented by the Spring Research and Manufacturers' Association (SRAMA) in collaboration with Warwick University's Interdisciplinary Mathematical Research Programme (IMRP). Both lie at the core of multi-million pound industries.

FRACTAL IMAGE COMPRESSION

When you buy a computer nowadays you generally get a lot of free 'bundled' software. Among the demos and games you often find a complete encyclopaedia on a single CD-ROM. You may, in passing, marvel at the technological advances required to cram so much information into so small a space. But there's more to it than just technology. There's new mathematics, too, which makes it possible to cut the amount of information down to a manageable size before the technology even begins to cram it in.

Storing text is easy: its images that cause the trouble. The Encarta™ CD-ROM encyclopaedia, for instance, includes eight thousand colour pictures. The information in a picture is can take up as much space as a hundred pages of text. A screen image is made up from millions of tiny 'pixels' — individual dots of colour. Today's technology cannot store eight thousand images on one CD-ROM if they are represented in this pixel-based fashion. But there they are, in exquisite detail, in the Encarta™ encyclopaedia.

The information in a page of text can be reduced considerably by taking advantage of the redundancy of natural language, which represents information rather inefficiently. But how do you compress a picture? The most important message from fractals is that simple processes can produce complex forms. Michael Barnsley, a British-born mathematician living in the US, realised that this peculiarity of fractals might form the basis of a method of image compression. The key example that drove his thoughts was a fern. A fern consists of a large number of fronds arranged two by two along a stem. Moreover, each frond is like a miniature fern, having its own frondlets, and so on. The entire structure of the fern can therefore be represented by a 'collage' in which the original fern is dissected into four 'transformed' — distorted — copies of itself. These copies are the two largest fronds, near the base of the stem; the rest of the fern above those fronds; and the tiny bit of stem left at the base — which, in a bit of a cheat, is thought of as a fern squashed flat.

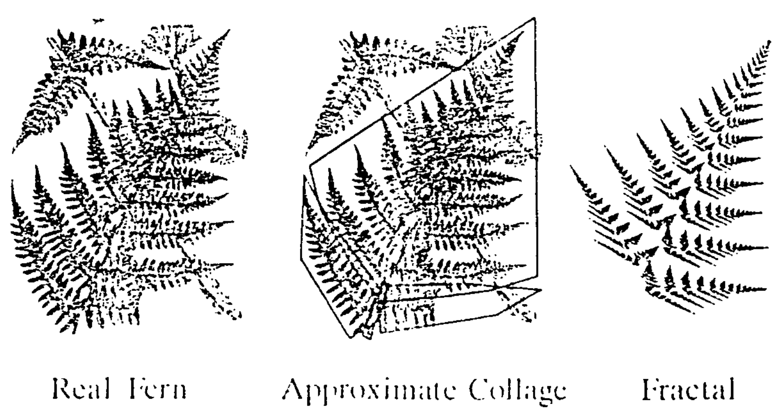

Fig.2 Barnsley's fern.

Like the old Chinese recipe that begins 'boil twenty chickens in a pot, then throw away the chickens,' the trick now is to throw away the fern. Barnsley found a method to regenerate the entire fern, in full detail, from the set of mathematical transformations (known as an **Iterated Function Scheme** or **IFS**) that determines the collage. One way is repeatedly to pick transformations at random from the set, and apply them to create a sequence of points in the plane. This sequence converges to the desired shape with probability 1 — that is, in practice, always.

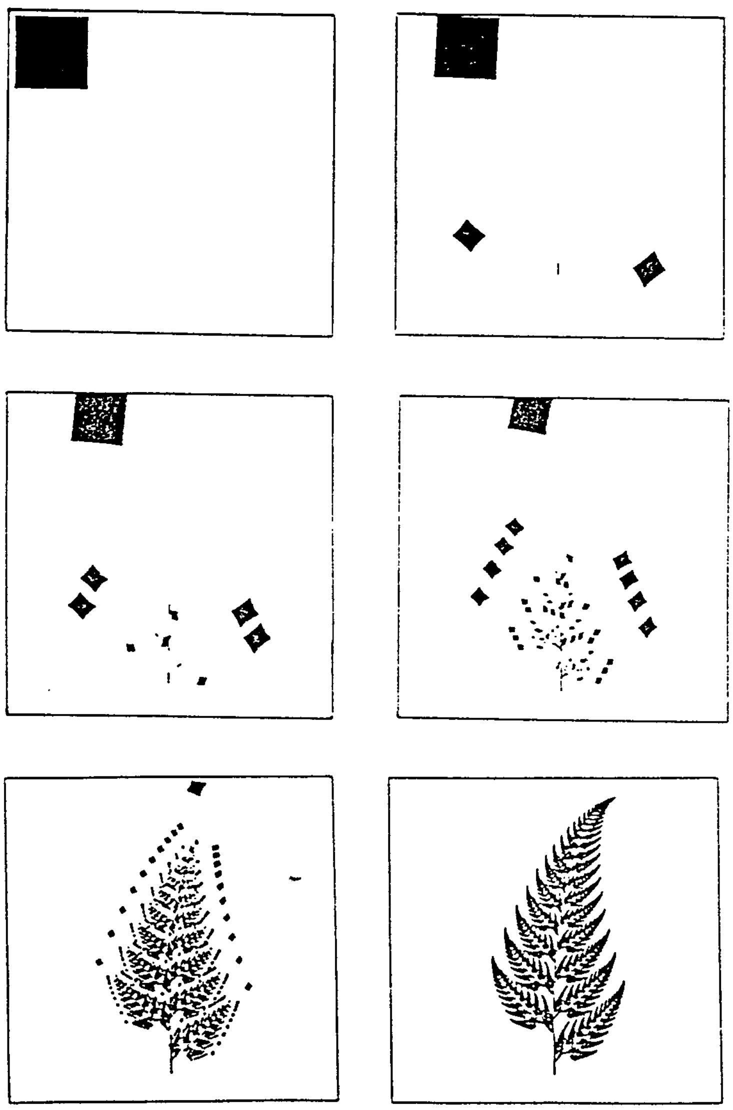

Fig.3 Reconstructing a shape from an IFS.

Each transformation can be specified by just six numbers, so twenty-four numbers suffice to generate the incredibly intricate shape of the fern. And the computer disk need store only those twenty-four numbers, not the million needed to determine the shape one tedious pixel at a time. At this point Barnsely tried to interest commercial companies in the potential of his scheme. No takers: the idea was just too original. So he formed his own company to develop the idea.

Ferns are all very well, but what about completely general pictures? Can they be similarly broken down into distorted copies? The answer is 'yes' — but you don't use copies of the entire picture. What you do is find tiny bits of the picture that resemble other, larger bits. Given enough such pieces, you can again represent the image in terms of a collage of mathematical distortions — and the list of numbers needed to specify them is far, far shorter than a pixel-by-pixel description. With Barnsley's method an image can be compressed to take up less than a hundredth of the usual space.

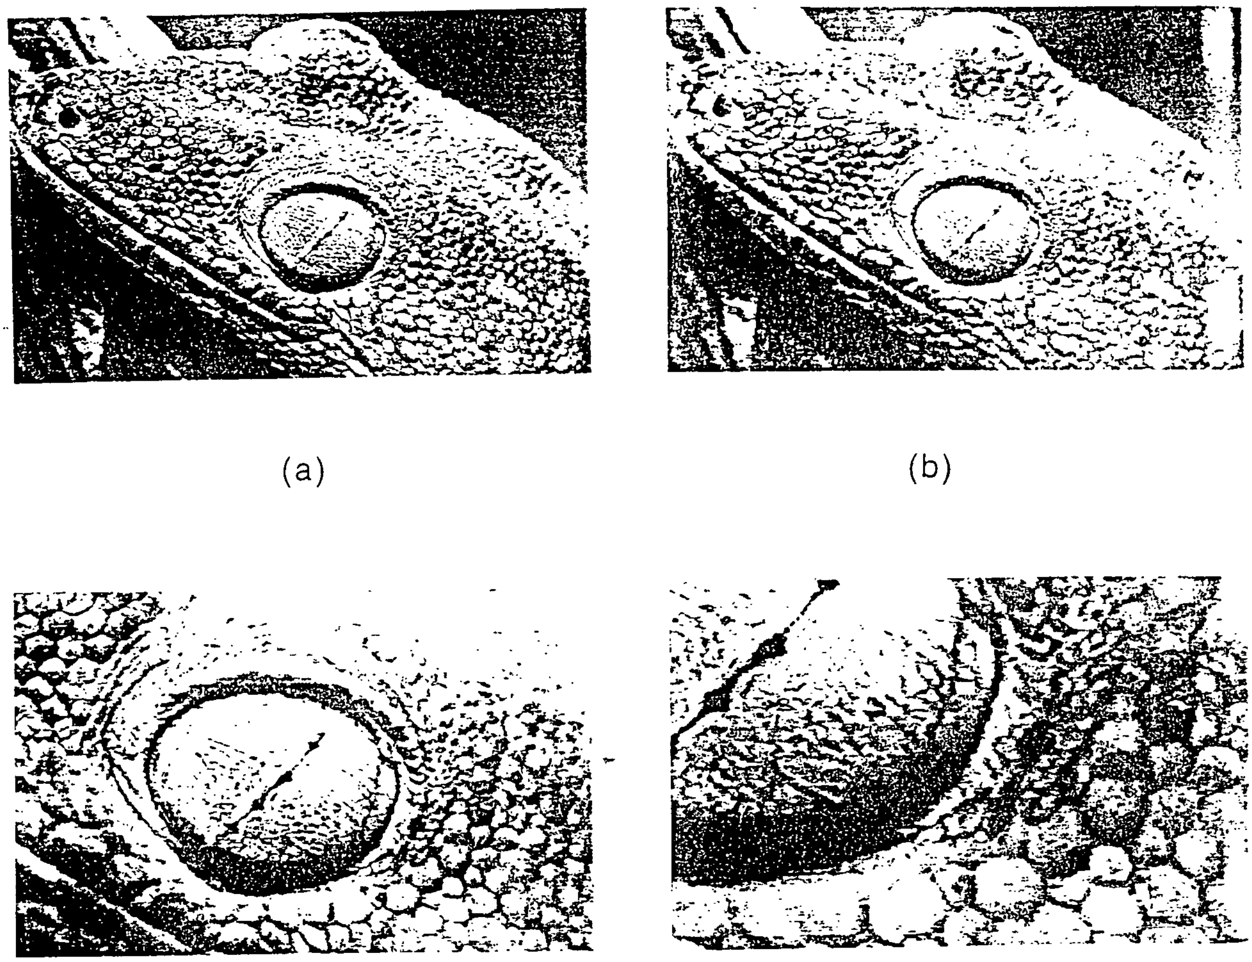

(a) (b) (c) (d)

Fig.4 Compression of a real image. The original image (a) of the gecko, a 640x400 resolution 24-bit image, requires 768 Kb of storage. (b) Compression ratio 156:1. (c) Compression ratio 625:1. (d) Compression ratio 2500:1.

The idea that was too wild for any commercial company to invest in it is now the basis of a multi-million dollar company, Iterated Systems Inc. Image-compression is important in many other areas too: for example satellite TV. Barnsley's company has just announced a fractal-based method for compressing video images. If you can compress the information in a TV picture by 50% then you can get twice as many channels on the same satellite — saving hundreds of millions of pounds.

FRACMAT

The dynamical aspect of fractals is known as **chaos.** The application of fractals and chaos that I want to tell you about is a machine called FRACMAT. FRACMAT applies fractal techniques to solve a problem that has plagued Britain's springmaking industry, and the wire industry that serves it, for at least twenty-five years. That problem is: how can you tell, quickly and cheaply, whether a consignment of wire can be coiled successfully into springs?

Springs are made by feeding wire, at speed, into a coiling machine, which is about the size and shape of two filing cabinets. The wire begins as a loose coil about a metre across, which lies horizontally on a rotating turntable. The coiling machine draws it along through a series of rollers, and runs it past two tools. One bends the wire through a quarter of a circle, and the other nudges it sideways. A third cutting tool snips off the spring when it has fully formed. A normal coiling machine can make three or four springs every second, although some specialist machines that use fine wire can coil 80,000 high-precision springs an hour — that's 22 per second.

Springs may look crude, but they are very precise components, and they have to be the right size, shape, and strength. There are springs everywhere. A video-recorder contains hundreds. Car engines contain between eight and thirty-two valve springs, depending on design. The collision-detecting device that sets off a car's airbag, to save the driver's face from intimate contact with the steering-wheel, is basically a ball balanced on a few springs. You wouldn't be too pleased if your airbag went off by mistake, so those springs have to be very precise, allowing the device to distinguish reliably between hitting an obstacle and driving over your local council's traffic-calming measures. Nowadays manufacturers are starting to put airbags in off-the-road vehicles, whose normal mode of operation is only marginally distinguishable from colliding with an avalanche. The airbag must be triggered only when the vehicle hits a rock big enough to stop it. So those springs have to be very precise indeed.

All of which poses a quality control problem.

It takes a skilled operator a lot of time to 'set' a coiling machine — that is, to adjust it so that it produces the right design of spring. Four to six hours is not unusual, and that's using computerised adjusters instead of spanners as in the Old Days. And that's also assuming that the wire has 'good coilability' — which means that if you adjust the machine right, it actually will form into springs. What the operator does is to coil some test springs, run them through the entire manufacturing process — including, if appropriate, heat treatment, hardening, galvanizing, and machining (say to get the ends flat). Then the resulting batch of springs is given a statistical test to see if it's of acceptable quality. If not, the operator tries to guess what went wrong, resets the coiling machine, and has another go... If the wire won't coil, this is never going to work, but the operator may take 12 hours or more before becoming convinced that the wire is at fault.

Prior to FRACMAT there was no quick and easy way to distinguish good coilability wire from poor coilability wire. All the wire sent to the springmakers by the wiremakers passes all of the standard quality control tests, things like material composition and tensile strength. Even so, about ten per cent of wire that passes these tests has poor coilability.

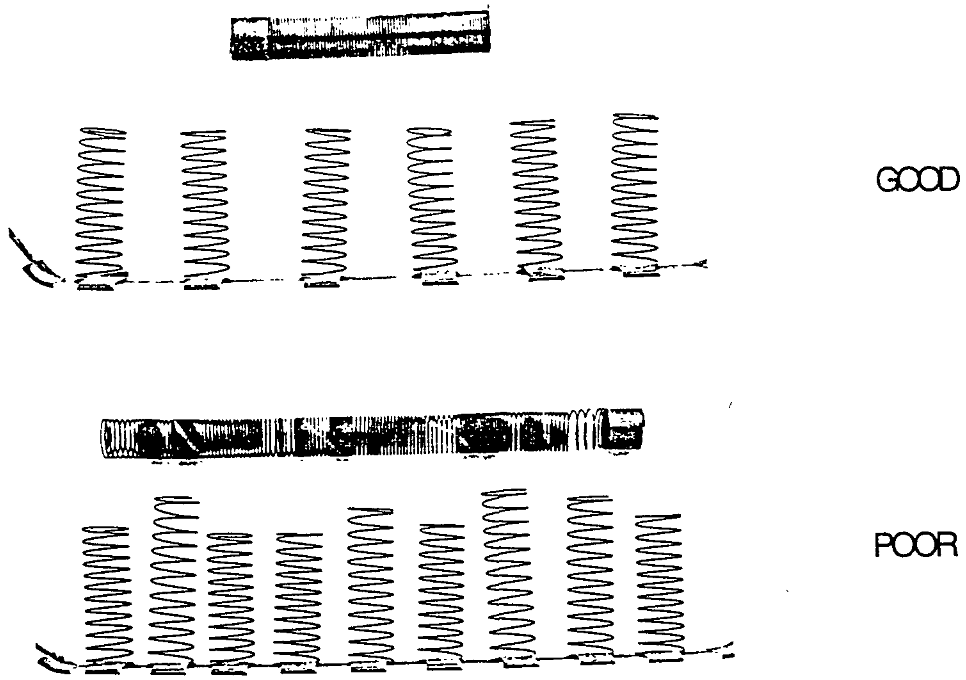

Fig.5 Wire with good and poor coilability.

Worldwide, the manufacture of springs is carried out by relatively small companies, and the same goes for the wiremaking companies that supply them with raw materials. These small companies have banded together to create their own joint R&D arm: SRAMA. SRAMA had devised what ought to be an effective text for coilability: force the wire sample to form a long spring by winding it round a long metal rod, or mandrel. Having done so, you can try to decide whether the test spring looks like what you'd expect from good coilability wire. There are several ways to do this. One is to get an experienced engineer and show him the spring: he either nods or shakes his head. It works, but you can't turn it into an international quality standard.

Another is to measure the spacings between successive coils on the test spring. Experiments show that on the whole, good coilability wire makes a test spring with nice, regular coils, whereas poor coilability wire makes a test spring with errratic spacings. SRAMA had invented a machine that did this with a laser micrometer, and fed the resulting list of numbers into a computer. They had then tried every statistical test in the book, and a few not in the book, trying to separate the good wire from the bad.

Nothing worked. It became clear that it isn't just the *statistics* of the spacings that matters, but the order in which they come. Say that a single coil is 'fat' if it's a bit wider than it ought to be, and 'thin' if it's a bit too narrow. Then, simplifying hugely, wire that produces successive coils something like

fat / thin / thin / fat / fat / thin / fat / thin / fat / thin / thin / fat / thin

will probably have good coilability, whereas wire that goes

fat / fat / fat / fat / fat / thin / thin / thin / thin / fat / thin / thin / thin

won't. The reason is that a real spring consists of several coils. In the first case, the errors tend to cancel out. In the second case, what you get is a complete spring that is too long, followed by one that is too short, and neither are any use.

This is an oversimplification, of course, but the basic idea is right. Fig.5 shows measured sequences of spacings for some typical wire samples, one very good, the other very bad. You don't have to be a genius to tell the difference; but it's in the no–man's-land between that the problem lies.

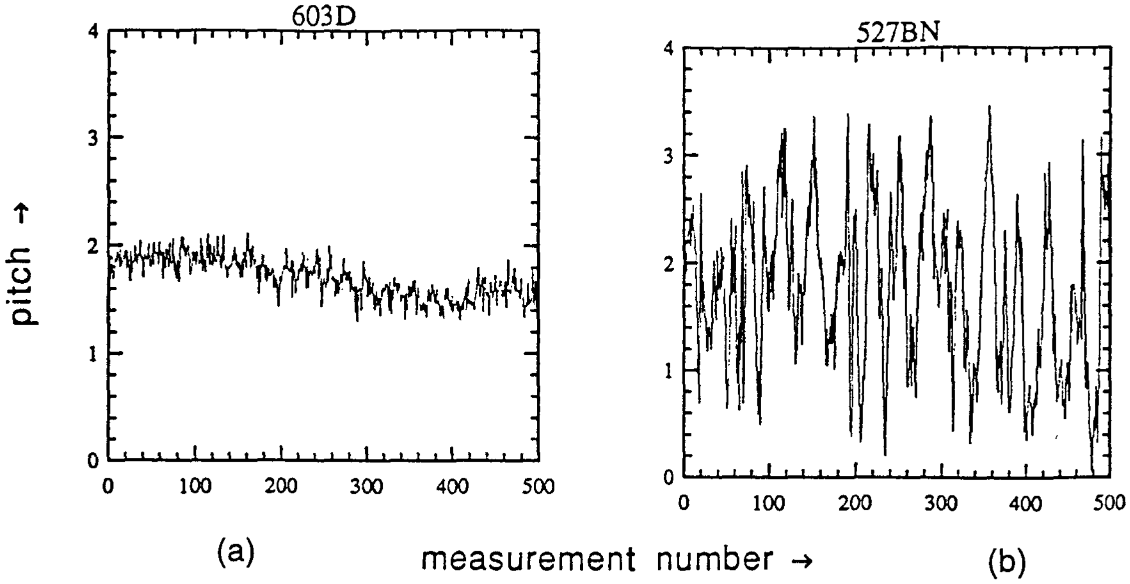

Fig.6 Sequences of spacings for some typical wire samples (a) good, (b) bad.

The key to coilability is *sequential* variability of the material properties of the wire, not statistical variability. But how can you quantify this? Chaos Theory, which is intimately entwined with fractals, provides a technique called phase space reconstruction, in which a time series of measurements is turned into a geometrical shape, called an *attractor*, by an appropriate mathematical algorithm. For chaotic dynamics, this attractor is a fractal, hence the connection. The sequence of coil-spacings produced by SRAMA's laser micrometer is, in effect, a time series. Phase space reconstruction makes the difference between good and bad coilability wire obvious and quantifiable. The reconstructed attractor generally resembled an elliptical blob. If that blob is nice and compact, the wire is good; if not, then the wire is bad.

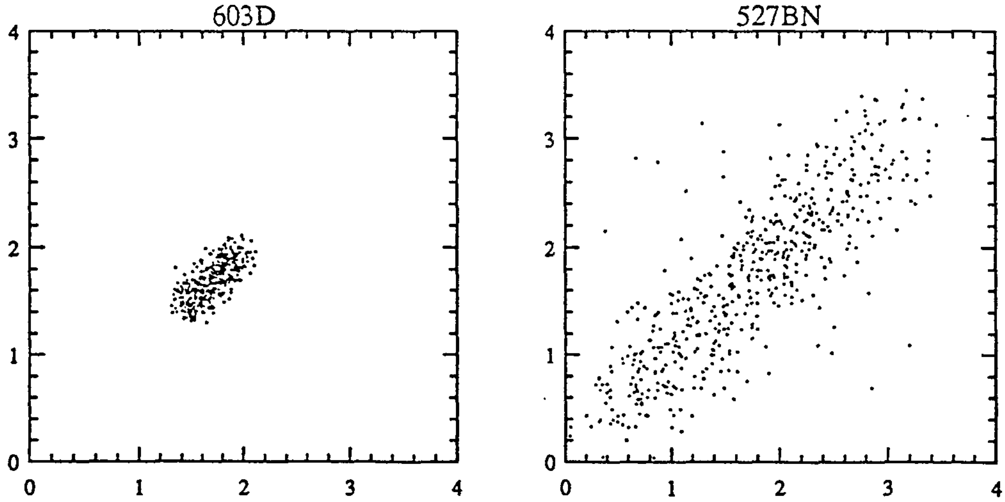

Fig.7 Reconstructed atractors for the two wires of Fig.6.

These 'attractors' don't *look* fractal, like the Mandelbrot set: they're just fuzzy clumps. This is either because there is a lot of random 'noise' in the manufacturing processes, or because most higher-dimensional attractors look like fuzzy clumps anyway. Indeed these may be two ways to say the same thing. However, it doesn't *matter* whether the attractor for a test spring is genuinely fractal or not. The important practical question is to find a quantitative method for distinguishing good wire from bad — and to say *how* good or bad it is. Phase space reconstruction doesn't just work for a chaotic time series. It works for *any time series whatsoever*; and it provides a rigorous mathematical way to characterise the type of sequential variability occurring in that time series.

Funding for the FRACMAT prototype and software development was supplied by the Department of Trade and Industry's 'Carrier Technology' programme. A team at SRAMA (Len Reynolds, Derek Saynor, Mike Bayliss, and others) joined forces with a team at Warwick University (myself, Mark Muldoon, Matt Nicol). They were supported by several wire manufacturers who supplied samples of wire, told them which they thought were good or bad, tested out protoptype equipment in their factories, and attended regular project meetings. A few key people from the DTI kept a close and helpful eye on the project. Over a project period of just two years, SRAMA designed and built a machine that would automatically form a test spring (with about 500 individual coils) on a mandrel, and Warwick embodied various chaos-theoretic algorithms for phase space reconstruction, together with some more traditional algorithms for detecting periodic variations, in a computer program. It was suggested that we call the machine MANDRELBOT, but we chickened out and named it FRACMAT — the acronym stands for FRACtal MATerials

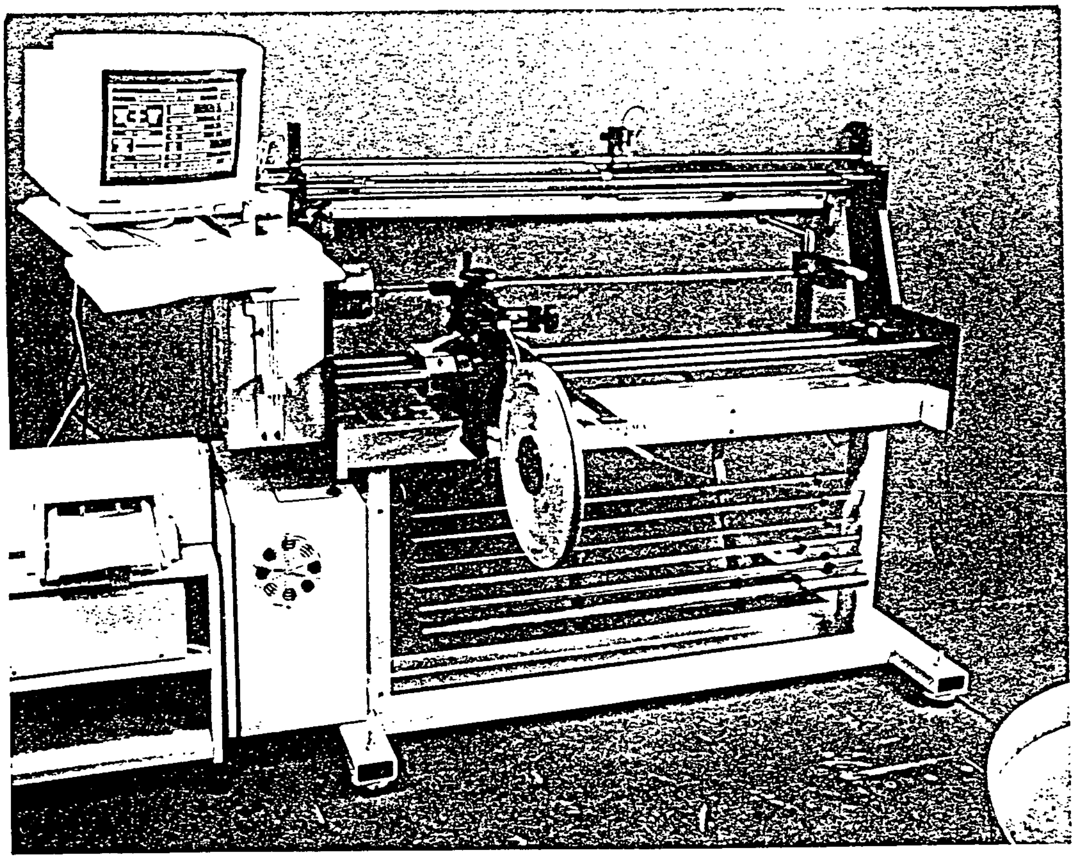

Fig.8 The FRACMAT machine

FRACMAT is about the size of a large desk turned on its side. Its computer both controls the operation of the machine and analyses the results. It has two motors: one turns the mandrel about which the test spring will be coiled, and the other, via a worm drive, moves the wire along the mandrel so that it is always being coiled in the right place. The computer counts the number of turns and stops when a required number — usually 500 — is reached. The completed spring is transferred (by hand, and with a few jiggles to make sure it settles freely) to a shelf above which the laser micrometer can track rapidly along, measuring all the spacings. Virtually instantly, the machine reconstructs the attractor, quantifies how long and how broad it is, and plots out the results on a 'classification diagram that determines how good or bad the wire's coilability is. The entire test takes about three minutes — not so long ago it took SRAMA two days to make the same measurements.

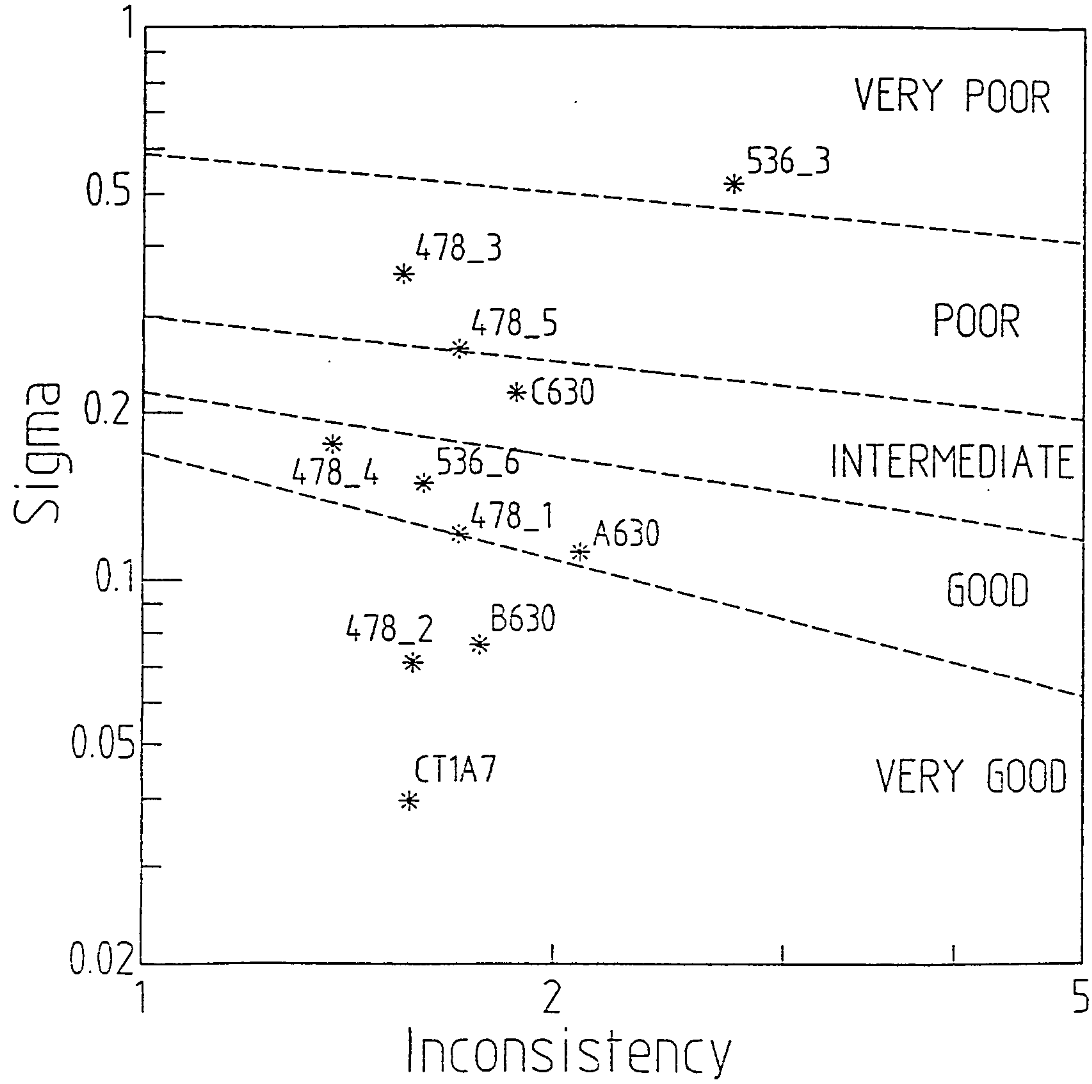

Fig.9 The classification diagram.

FRACMAT is now the subject of a UK patent application. It won joint second prize in the 'innovative metrology' section of the Metrology for World Class Manufacturing Awards in Birmingham in 1995. If it is adopted as widely as anticipated, it should save the UK springmaking sector tens of millions of pounds very year.

Further Reading

FRACTAL IMAGE COMPRESSION

Michael F Barnsley and Lyman P Hurd, *Fractal Image Compression*, A K Peters 1993.

Michael F Barnsley, *Fractals Everywhere* (2nd ed.), Academic Press 1993.

FRACMAT

M.Muldoon, M.Nicol, L.Reynolds, and I.Stewart, Chaos Theory in quality control of spring wire — Part I, *Wire Industry* **62** (1995) 309-311; Part II, *Wire Industry* **62** (1995) 491-492; Part III, *Wire Industry* **62** (1995) 492-495.

M.Bayliss, M.Muldoon, M.Nicol, L.Reynolds, and I.Stewart, The FRACMAT test for wire coilability: a new concept in wire testing, *Wire Industry* **62** (1995) 669-674.

GRESHAM COLLEGE

Policy & Objectives

An independently funded educational institution, Gresham College exists

- to continue the free public lectures which have been given for 400 years, and to reinterpret the 'new learning' of Sir Thomas Gresham's day in contemporary terms;
- to engage in study, teaching and research, particularly in those disciplines represented by the Gresham Professors;
- to foster academic consideration of contemporary problems;
- to challenge those who live or work in the City of London to engage in intellectual debate on those subjects in which the City has a proper concern; and to provide a window on the City for learned societies, both national and international.

Founded 1597

Gresham College, Barnard's Inn Hall, Holborn, London EC1N 2HH
Tel: 020 7831 0575 Fax: 020 7831 5208
e-mail: enquiries@gresham.ac.uk
Web site : www.gresham.ac.uk