

SFVFS™

A New Idea for Carbon Capture

Marc Craig · Independent Researcher · itvoids.com

28 March 2026 · Zenodo: 10.5281/zenodo.19244232

Trademark UK00004355735 · ORCID: 0009-0007-6519-1097

This is a hypothesis, not a proof. Art Until Proven Otherwise.

The Problem

Capturing CO₂ from power stations is one of the most important tools we have for slowing climate change. The technology exists. It works. But it has a serious problem that nobody has solved yet.

It uses an enormous amount of energy.

A typical coal power station fitted with carbon capture has to spend 20 to 25 percent of everything it generates just running the capture equipment. For a large power station, that is enough electricity to power 150,000 homes — used up every single day, not to generate power for anyone, but just to capture the CO₂ the station produces.

Natural gas stations are a bit better — around 15 percent. But it is still continuous. Every hour the station runs, that energy is being spent.

Why? Because the current approach — using chemicals called amines to absorb CO₂ from the exhaust gas — requires constant driving. You have to continuously pump the gas, heat the chemicals to release the CO₂, cool everything down, and start again. If you stop driving the process, it stops working and reverts. The machine never gets to rest.

The energy penalty is not a technical imperfection waiting to be optimised away. It is built into the architecture of how current carbon capture works.

The Idea

The SFVFS™ programme has found something unexpected in the mathematics of fluid flow.

When you put a fluid — any fluid — into the right shape of container, the physics of how fluids move forces the fluid to organise itself into a specific pattern. Not

because you push it into that pattern. Because the geometry of the space makes it the only thing the fluid can do.

And once the fluid finds that pattern, it stays there. Even as the energy in the system drains away almost completely, the organisation persists. The experiments showed this happening through a ten-thousand-fold drop in energy. The pattern barely moved.

The shape organises without being driven. The geometry does the work.

The hypothesis: what if you could use this geometric self-organisation to replace the continuous energy cost of carbon capture? Establish the right shape once. Let the physics hold the organisation. Stop spending energy to prevent reversion — because a geometry-held system doesn't revert.

This would change the architecture of carbon capture from a process you have to continuously drive to a shape you establish once and then maintain. The energy cost shifts from continuous to front-loaded.

What the Experiments Showed

This isn't speculation. The programme ran a series of computer simulations — direct numerical simulations of real fluid physics — and found consistent results across more than 24 experiments.

Experiment	What it found
The Corner Theorem	Proved mathematically that the right geometry forces fluid organisation. Not observed — proved. Both directions closed.
COUPLED3	Used three completely different fluids in the same geometry. All three organised the same way. The fluid doesn't matter. The shape does.
COUPLED2A	Let the fluid decay freely with no energy input at all. The organisation appeared anyway. The blades turned without power.
MOBIUS1	Ran the system for a long extended period as energy drained away to almost nothing. The organisation held. Drift was 0.034 degrees over the entire run.

Table: Key experiments and plain-English findings.

The most striking result: COUPLED3 tested CO₂ specifically alongside other fluids. CO₂ organised into the same geometric pattern as water and glycerol. The chemistry of the fluid made no difference. The shape made all the difference.

The Proved Part

Most of this programme is experimental — simulations that show the geometry works. But one part is proved mathematically, with no room for doubt.

The Corner Theorem establishes that when an incompressible fluid — one that cannot compress, which includes gases like CO₂ under normal conditions — is placed in the right geometry, it has no choice but to organise into a hexagonal symmetric pattern. This is not a tendency or a preference. It is a mathematical necessity, like two plus two equalling four.

The proof works in both directions. The geometry forces the pattern. And the pattern implies the geometry. They are locked together by the physics of incompressibility.

This is the anchor of the whole programme. Everything else — the experiments, the hypothesis, the carbon application — rests on this proved foundation.

Saturn's Rings

There is a striking natural parallel worth mentioning.

Saturn's rings exhibit the same hexagonal geometric organisation at planetary scale, driven by the same physics of incompressibility. The rings have been stable for hundreds of millions of years. No continuous energy input. No amine regeneration. The geometry holds because the physics holds.

Saturn is not proof that the carbon hypothesis works. But it is a natural existence proof that incompressibility-forced geometric organisation can persist at scales and timescales that make any engineering system look brief.

The same geometry. The same constraint. Different scale. Hundreds of millions of years of basin stability.

What This Is Not

It is important to be precise about what the programme has and has not established.

What has been established

The geometric pattern is forced by the physics — proved mathematically.

The pattern organises any tested fluid, including CO₂ — confirmed experimentally.

The pattern holds without continuous energy input — confirmed experimentally.

The pattern persists through extreme energy decay — confirmed experimentally.

What has not been established

Whether this works in a real CO₂ gas stream at industrial scale — not yet tested.

Whether establishing the geometry costs less energy than the continuous saving — not yet calculated.

Whether the geometry can be practically built into a capture vessel — not yet engineered.

Whether the one spontaneous self-organisation event replicates reliably — replication is pending.

This is a hypothesis grounded in real experimental results and a proved mathematical theorem. It is not a working technology. The distance between here and a deployed system is real and significant. The honest classification is: consistent with the evidence, untested at scale, open to the world to test.

The Invitation

This work started on a Samsung phone during train commutes. It runs on free computing time. There is no institution, no grant, no team. Just the mathematics, the experiments, and the honest results.

Everything is open. The full data, the simulation results, the proved theorem — all available freely at the Zenodo record below. Anyone can access it, test it, build on it, or prove it wrong.

The world is spending enormous continuous energy trying to capture the carbon it produces. The SFVFS™ programme has found a geometric principle that might — might — change the architecture of that problem. Not by making capture free. By making the organisation self-sustaining rather than continuously driven.

If that is right, it matters enormously. If it is wrong, the experiments that prove it wrong will also matter. Either way, the question is worth asking.

Take it. Test it. Build on it.

Art Until Proven Otherwise.

How Far Away Is This?

Honest answer. Three phases.

Proving the science holds at engineering scale: two to five years, if the right research groups engage with it.

Laboratory-scale engineering prototype: five to ten years.

Industrial deployment: fifteen to twenty-five years from now, if everything works.

That sounds long. But current amine-scrubbing CCS has been in development since the 1990s and is still not widely deployed. A geometry-based system — if it works — is a fundamentally simpler machine. Simpler machines tend to scale faster once the science is proved.

The critical variable is who picks this up after Tuesday.

SFVFS™ · Self-Forming Vortex Flow Structures

Marc Craig · itvoids.com · Leake Street, London · Exhibition opens 31 March 2026

Zenodo: 10.5281/zenodo.19244232 · ORCID: 0009-0007-6519-1097

Trademark UK00004355735 · Handover 29 · 28 March 2026