

Design of Programmable Matter

by

Ara N. Knaian

S.B. Electrical Science and Engineering (1999)
Massachusetts Institute of Technology, Cambridge, MA

M. Eng. Electrical Engineering and Computer Science (2000)
Massachusetts Institute of Technology, Cambridge, MA

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
in partial fulfillment of the requirements for the degree of

Master of Science in Media Arts and Sciences

at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2008

© 2008 Massachusetts Institute of Technology
All Rights Reserved

Author _____
Program in Media Arts and Sciences
January 14, 2008

Certified by _____
Neil A. Gershenfeld
Associate Professor of Media Arts and Sciences
Director, Center for Bits and Atoms

Accepted by _____
Deb K. Roy
Associate Professor of Media Arts and Sciences
Chair, Departmental Committee on Graduate Students

Design of Programmable Matter

by
Ara N. Knaian

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning
on January 14, 2008
in partial fulfillment of the requirements for the degree of
Master of Science in Media Arts and Sciences

ABSTRACT

Programmable matter is a proposed digital material having computation, sensing, actuation, and display as continuous properties active over its whole extent. Programmable matter would have many exciting applications, like paintable displays, shape-changing robots and tools, rapid prototyping, and sculpture-based haptic interfaces. Programmable matter would be composed of millimeter-scale autonomous microsystem particles, without internal moving parts, bound by electromagnetic forces or an adhesive binder.

Particles can dissipate 10 mW heat, and store 6 J energy in an internal zinc-air battery. Photovoltaic cells provide 300 μ W outdoors and 3.0 μ W indoors. Painted systems can store battery reactants in the paint binder; 6 J / mm³ can be stored, and diffusion is fast enough to transport reactants to the particles. Capacitive power transfer is an efficient method to transfer power to sparse, randomly placed particles. Power from capacitive transfer is proportional to V_{DD}^2 : 100 μ W at 3.3V and 12 mW at 35V. Inter-particle communication is possible via optical, near-field, and far-field electromagnetic systems. Optical systems allow communication with low area (sub-mm) particles, and 24 pJ/bit. Near-field electromagnetic gives precisely controlled neighborhoods, localization capability, and 37 pJ/bit. Far-field radio communication between widely spaced particles may be possible at 60 GHz; antennas that fit inside 1 mm³ exist; complete transceivers do not. A 32-bit CPU uses less than 0.26 mm² die area, 256K x 8 SRAM uses 1.1 mm², and 256K x 8 FLASH uses 0.32 mm². Direct-drive electric and magnetic field systems allow actuation without moving parts inside the particles. Magnetic surface-drive motors designed for operation without bearings are not power-efficient, and parasitic interactions between permanent magnets may limit their usefulness at millimeter particle dimensions. Electrostatic surface-drive motors are power-efficient, but practical only at particle dimensions below a few millimeters.

We constructed a prototype paintable display; a distributed PostScript rendering system with 1000 randomly-placed 3.4 cm nodes, each with a CPU, IR communications, and LED. The system is used to render the letter "A." We present a design, not yet constructed, for a literal paintable display, with 1.0 mm rendering particles, each with a microprocessor and memory, and 110 μ m display particles, with tri-color LED's and simpler circuitry. Storage of zinc-air battery reactants in the paint binder would provide an 8 hour battery life, and capacitive power distribution would allow continuous operation.

We constructed a prototype sliding-cube modular robot, with 3.4 cm nodes. The system uses magnetic surface-drive actuation. We demonstrate horizontal lattice-unit translation. We describe a design, not yet constructed, for a sliding-cube modular robot with 2 mm nodes. The cubes use standard-process CMOS IC's, inserted into a cubic space frame and wire-bonded together. Arrays of passivated electrodes, 1 μ m from the surface of the cubes, are used for electrostatic surface-drive actuation, zero-power latching, power transfer, localization, and communication. The design allows actuation from any contacting position. Energy is stored in a standard SMT capacitor inside each node, which is recharged by power transfer through chains of contacting nodes.

Thesis Supervisor: Dr. Neil A. Gershenfeld
Title: Associate Professor of Media Arts and Sciences
Director, Center for Bits and Atoms

Design of Programmable Matter

by

Ara N. Knaian

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning
on January 14, 2008
in partial fulfillment of the requirements for the degree of
Master of Science in Media Arts and Sciences

Thesis Supervisor _____
Neil A. Gershenfeld
Associate Professor of Media Arts and Sciences
Director, Center for Bits and Atoms

Thesis Reader _____
Joseph M. Jacobson
Associate Professor of Media Arts and Sciences
MIT Media Lab

Thesis Reader _____
Daniela L. Rus
Professor of Electrical Engineering and Computer Science
Co-Director, CSAIL Center for Robotics

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGEMENTS

This work was supported in part by the Center for Bits and Atoms, NSF, DARPA, and DTO.

Thanks first to Neil Gershenfeld, for creating such a spectacularly interesting place to work, and assembling such a wonderful group of people with whom to do that work. Thanks to Joe Jacobson and Bill Butera, for the free-wheeling discussions that posed the questions that this thesis attempts to answer. Without them, this thesis would probably have been about some other topic entirely. Thanks to Neil Gershenfeld, Joe Jacobson, Daniela Rus, Carol Livermore, Saul Griffith, Mark Pavicic, and Bill Butera, for being my mentors: pointing out literature I needed to read, helping me sharpen my ideas, and through their feedback and guidance making my work so much better.

Very direct thanks go to Forrest Green, David Greenspan, and Monica Sun, exceptionally talented undergraduate researchers with whom it was a pleasure to work, and who will certainly recognize many of the ideas in this thesis as springing from one of our late-night conversations. Beyond this, I want to point out that David and Monica wrote 100% of the code for the process-fragment paintable display, a huge job.

Thanks to my lab-mates and fellow students; although we work on individual projects in theory, all of our work bears each other's undeniable imprint in practice. In particular, I want to thank Amy Sun, Nels Peterson, Kenny Cheung, Manu Prakash, George Popescu, David Darlyrymple, Jon Santiago, David Kopp, Kerry Lynn, Ilan Moyer, Luis Lafuente, Kailang Chen, Mike Hennebry, Marty Vona, Dan Goldwater, Jason Taylor, Ben Recht, Xu Sun, Raffi Krikorian, Yael Maguire, Ben Vigoda, Brian Chow, Jaebum Joo, and David Mosley. I can remember several times when there was a deadline to meet, and everyone around the lab pitched in to help out. I can vividly remember Amy and Kenny doing lasercutter magic to build cube frames to my absurd specifications, Nels prototyping magnet patterns, George assembling cubes, David and Luis helping me to debug the electrostatic tiles in the Fluorinert caves of Fargo, and Jason writing a script to automate programming 1000 processors.

Susan Bottari, Sherry Lassiter, Kelly Maenpaa, Mike Houlihan, and Joe Murphy did a spectacular job keeping supplies and equipment flowing into the lab at a sometimes breakneck pace, among plenty of other responsibilities. Thanks also to Linda Peterson, Gigi Shafer, and Marilyn Pierce, for making sure I had the right paperwork submitted on time, and generally looking out for me academically. John DiFrancesco and Tom Lutz run a great machine shop, and were always willing to lend a hand to help figure out how to get something built.

Thanks to Harry Keller, David Newburg, Mr. Martins, Mr. Wells, and Dr. Duffy.

Thanks to my parents, for the obvious, but also for getting me interested in science and engineering in the first place.

Finally, thanks to Linda, for her constant friendship, support, and love.

TABLE OF CONTENTS

	<u>Page</u>
1 INTRODUCTION	8
1.1 Personal Fabrication and Digital Materials	8
1.2 Programmable Matter	11
1.3 The Present Moment in Semiconductor Manufacturing	12
1.4 Scope	13
1.5 Background	14
2 PHYSICS, MATERIALS, AND DEVICES	18
2.1 Heat.....	19
2.1.1 Heat Transfer and Thermal Resistance.....	20
2.1.2 Thermal Capacitance and Thermal Time Constant.....	23
2.1.3 Pulsed Power Safe Operating Area Curves	25
2.2 Power	27
2.2.1 Particle-Stored Energy.....	28
2.2.1.1 Electrochemical Cells.....	28
2.2.1.2 Microengines.....	29
2.2.1.3 Radioisotopes.....	30
2.2.1.3.1 Heat Production	30
2.2.1.3.2 Toxicity	31
2.2.2 Binder-Stored Energy	32
2.2.2.1 Inside-Out Zinc Air Battery	32
2.2.2.2 Binder-Stored Fuel	34
2.2.3 External Power.....	34
2.2.3.1 Photovoltaic Cells.....	35
2.2.3.2 Reactive Power Transfer.....	35
2.2.3.3 Power and Efficiency Calculations	39
2.3 Communications.....	42
2.3.1 Introduction	42
2.3.2 Communications Transports.....	43
2.3.2.1 Optical	43
2.3.2.2 Near-Field RF Communication.....	44
2.3.2.2.1 Near-Field Inductive Communications System Design	44
2.3.2.3 Propagating-Wave RF Communication	49
2.4 CPU and Memory.....	50
2.4.1 CPU.....	50
2.4.2 RAM	50
2.4.3 FLASH	51
2.5 Light Emitting Diodes	51
2.5.1 Power Requirements	51
2.6 Actuation	53
2.6.1 Why Distributed Action?	53
2.6.2 Electrostatic vs. Magnetic	54
2.6.3 Magnetic Field Systems.....	54
2.6.3.1 The case against magnetically actuated “Utility Fog”	54
2.6.3.2 Permanent-Magnet Linear Motors	59
2.6.3.3 Scaling Analysis of Permanent-Magnet Interactions	64
2.6.4 Electric-Field Systems	65
2.6.4.1 The difficulty with Electrostatic “Utility Fog”	65
2.7 Integration of Dissimilar Process Technologies	68
2.8 Economics.....	68
3 APPLICATION EXAMPLES.....	70
3.1 Paintable Display.....	70
3.1.1 Centimeter-Scale Prototype.....	71

3.1.1.1	Materials and Methods.....	71
3.1.1.2	Results and Conclusion	73
3.1.2	Millimeter-Scale Design	74
3.1.2.1	Display Particles.....	74
3.1.2.2	Rendering Particles	76
3.1.2.3	Power	77
3.1.2.3.1	Power: Random Environment.....	77
3.1.2.3.2	Power: Controlled Environment.....	77
3.2	Modular Robot.....	78
3.2.1	Centimeter Scale Prototype: Magnetic Cubes.....	78
3.2.1.1	The Quandry of Magnet Placement	78
3.2.1.2	Geometric Shielding of Magnets using Nubs.....	80
3.2.1.3	Design of the Coil	82
3.2.1.4	Materials and Methods.....	84
3.2.1.5	Results: Volume and Area Usage of Subsystems	89
3.2.1.6	Fabrication Technology.....	90
3.2.1.7	Experimental Experience with the System.....	90
3.2.2	Centimeter Scale Prototype: Electrostatic Tiles	92
3.2.3	Millimeter Scale Design	96
3.2.3.1	Energy Storage	97
3.2.3.2	The Power Supply	97
3.2.3.3	The Frame.....	99
3.2.3.4	The Actuation Chips.....	99
3.2.3.5	Latching Force Calculation.....	103
3.2.3.6	Motive Force Calculation.....	104
3.2.3.7	Shaping the Force Distribution.....	105
3.2.3.8	Motion Speed and Power Consumption	106
3.2.3.9	Localization	107
3.2.3.10	Reconfiguration Geometry and Motion Planning	107
3.2.3.11	Capacitive Power Transfer	107
4	CONCLUSIONS.....	108
5	APPENDICES.....	109
5.1	The Magnetic Field of a Rectangular Permanent Magnet	109
5.1.1	Notes on Computation	113
6	REFERENCES	115

1 INTRODUCTION

1.1 Personal Fabrication and Digital Materials

Not too long ago, computers were specialized machines, usable only by experts, and affordable only by major corporations and universities. Today, personal computers are everywhere, usable by a five-year-old, and cost less than one month's rent. The on-going Personal Fabrication [1] revolution promises to do the same for fabrication tools; to give ordinary people the means to create high-technology objects at home.

Today's industrial manufacturing equipment is *analog* --- machine tools, for example, even if controlled by a digital computer, move a cutting tool over a real-valued path to shape an *analog material*. For mass production, where the tooling can be expensive but the cost per unit must be low, this is a fine approach.

By contrast, biological fabrication is *digital*. The ribosome assembles amino acids end-to-end, which fold into the proteins that result in all of the wonders of life. Amino acids are a digital material, because they have features that facilitate their own assembly according to a digital code. Another example of a digital material is LEGO, which can be easily assembled into discretized structures using the regular arrays of press-fit pins and sockets on each block.

Industrial manufacturing equipment is not something that most people would want in their homes, except maybe in the basement; by and large, it is messy, noisy, and dangerous. These characteristics are intimately linked to the fact that it is analog. Analog tools, additive or subtractive, work by performing operations that are extreme: impacting material with enough pressure to cause plastic deformation (machining), heating a solid material above its melting point (casting), or dissolving excess material with powerful solvents. (etching).

Unlike analog materials, digital materials are specifically designed to be easy to assemble --- so the processes used to make things from digital materials can be much friendlier, with no flood coolant, blast furnace, or fume hood required. With digital materials, of course, one is limited to making objects that can be assembled from blocks; but as we will see, the blocks can be very small, and the success of digital audio and digital printing have shown that with small enough blocks, the digital abstraction is not much of a limitation at all.

The ink-jet printer is a digital fabrication machine already present in many homes. An ink-jet printer prints arbitrary digital images by ejecting discrete droplets of specially designed ink onto paper. A three-dimensional ink-jet printer [2] fabricates three-dimensional objects by printing multiple layers. A multi-material three-dimensional printer that can print conductors, semiconductors, and insulators can print functional electromechanical devices. Fuller, Wilhelm, and Jacobson [3], have fabricated micro-actuators by ink-jet printing. (see Figure 1) Ridley, Nivi, and Jacobson [4] fabricated field-effect transistors, also by ink-jet printing. Malone and Lipson's

Fab@Home project [5] distributes free plans for a desktop multi-material 3D printer which can also embed packaged electronic components, and has been used to build a functioning LED flashlight. (Figure 2)

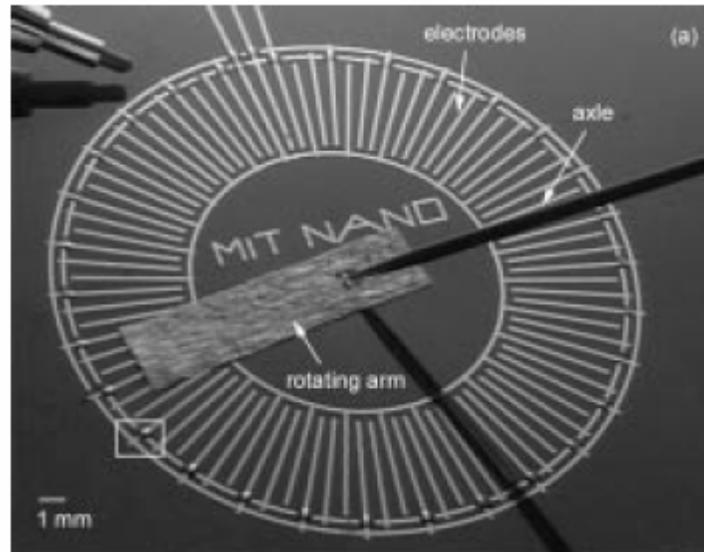


Figure 1: Ink-Jet printed electrostatic induction motor, from [3]

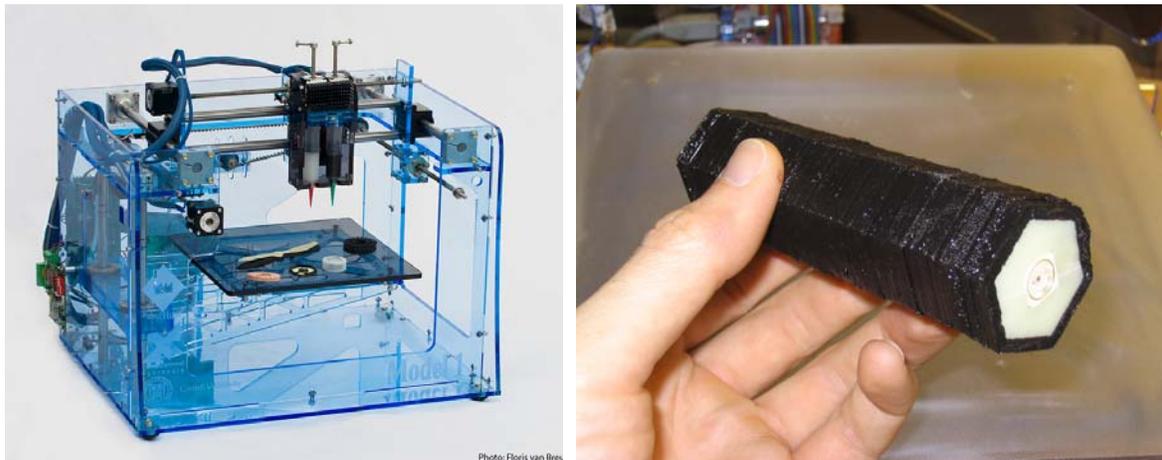


Figure 2: Left: Fab@Home Model 1, Right: Printed flashlight (from fabathome.org)

Another approach, *digital fabrication by folding*, works on the same principle as protein synthesis in molecular biology. Discrete building blocks are assembled into chains, in an order to code for the desired structure, and then the blocks fold themselves into that structure. Griffith [6] demonstrates a digital material consisting of square tiles with four different permanent magnet patterns and hinged links at the corners. To build an object, the tiles are assembled into a chain that codes for the structure, and then sequenced out of a tip, after which the magnets cause the chain to fold up into the desired shape. Rothmund [7] creates arbitrarily programmable shapes

from DNA molecules, by synthesizing a custom DNA strand to rater-fill the shape, and shorter “staple” strands to hold the shape together in two dimensions. The strands are then mixed together and assemble themselves into the desired shape. The resulting patterns have 6 nm pixels. (Figure 4)

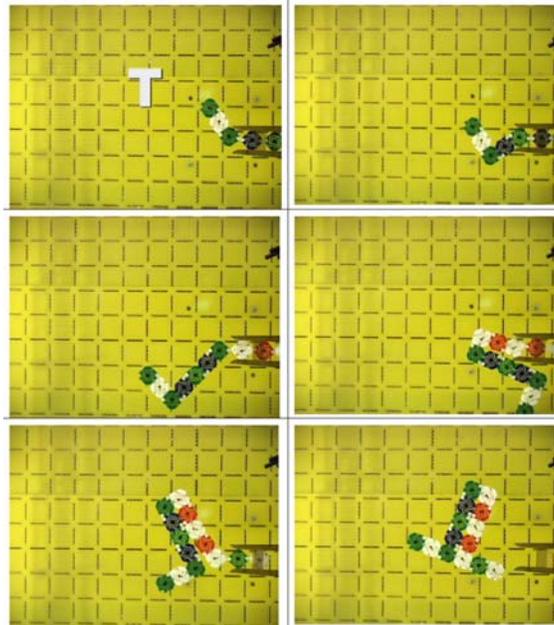


Figure 3: Digital fabrication of the letter “T,” by folding a sequence of magnetic tiles that code for the structure. From [6]

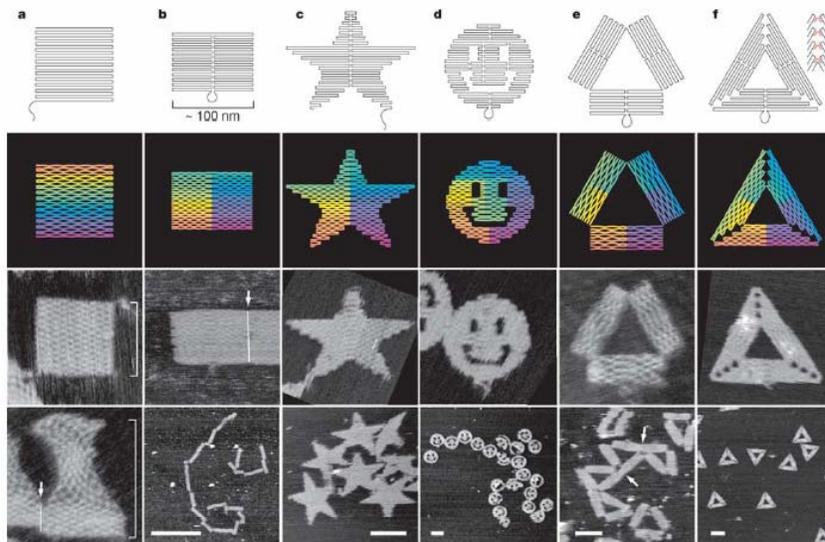


Figure 4: AFM Images of DNA Origami shapes, from [7]. The shapes are 100 nm across.

1.2 Programmable Matter

Digital materials can be categorized by the number of different kinds of blocks required, and the number of states each block can assume. [6] For example, biological protein synthesis uses 20 amino acid building blocks, each with essentially one state. Programmable matter is at the other end of the spectrum; only one kind of block is required, but the blocks can have an essentially infinite number of states, because each block contains a general-purpose computer.

Programmable matter is a proposed digital material having computation, sensing, actuation, and display as continuous properties active over its whole extent. A block of programmable matter would consist of thousands of engineered machines, which would work collectively to carry out the required functions. Programmable matter particles would be much more complex than the building blocks of other proposed digital materials, but batch photolithography processes make the cost of this additional complexity surprisingly low. In exchange for some additional material cost, no assembly tool at all is required; the material can assemble itself.

Like other digital fabrication technologies, programmable matter could be used as a material for fabricating static personalized objects --- say, a temperature-sensing coffee cup. But where programmable matter differentiates itself is in its ability to create transient objects, with lifetimes measured in seconds. Imagine an aircraft mechanic's tool, made from programmable matter blocks, that transforms itself from a wrench, to a screwdriver, to a walkie-talkie --- right in the mechanic's hand, while she is working. Another example --- a doctor could inject a patient with a few milliliters of programmable matter, and it could reconfigure into a submarine to travel to the site of a suspected problem, and then, if there was a problem, reconfigure into whatever surgical tool the doctor deemed necessary to solve the problem.

Programmable matter that was not biocompatible and could not bear huge mechanical loads would still enable a whole new class of human-computer interfaces --- for example, a haptic sculpture system, a shape-reconfigurable video-game controller, or a topography visualization system.

Even without actuation capability, programmable matter blocks with power, computation, communications, and light emitters, when mixed with a viscous liquid, could be packaged in cans as a paintable display. You could paint a quart of display paint onto your living-room wall and use it to watch television. If you wanted a larger screen, a higher resolution, or to widen your screen to support a new aspect ratio --- no problem --- just buy more display paint and paint it on.

At first blush, programmable matter might appear a far-off dream, bearing no relation to engineering reality. However, whether programmable matter is technologically feasible really boils down to one question: *Can one fabricate batches of small-enough machines that include*

sufficient computation, memory, power storage, power delivery, communication, sensing, and mutual actuation capability to accomplish the required task?

If millimeter-scale is small enough, then programmable matter may indeed be feasible: because a millimeter, although small by standards of human perception, is gigantic by the standards of the semiconductor industry.

1.3 The Present Moment in Semiconductor Manufacturing

The semiconductor industry has brought the manufacturing cost of a transistor below one ten-thousandth of a cent. This reality motivates the hypothesis that programmable matter could be manufactured using materials that are commercially available today, with only minor modifications to existing manufacturing processes.

In 1971, Intel Corporation released the 4004 microprocessor, which had four-bit data paths, a 108 kHz clock, and 2,250 transistors. Thirty-five years later, in 2006, Intel released the Core 2 Duo microprocessor, with two execution units, 32-bit data paths, a 3.0 GHz clock, and 291 million transistors. [8, 9, 10, 11]

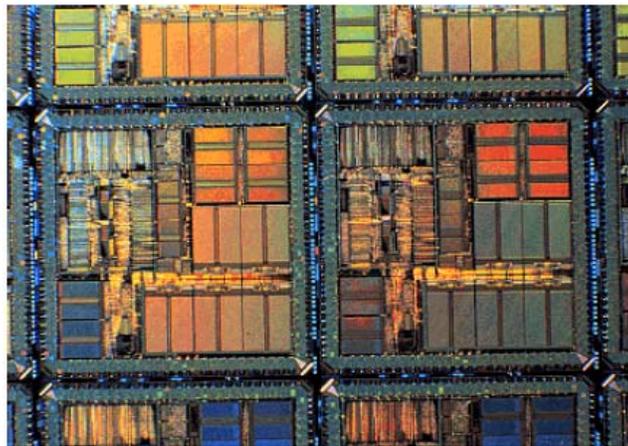


Figure 5: *Microprocessor Dice on a Silicon Wafer*

Imagine for the moment that you wanted to build a 4004 microprocessor today, but instead of using the original 10 μ m line-width process used by Intel, you used the 65 nm process. How much die area would you need?

$$4mm \times \frac{65nm}{10\mu m} = 26\mu m \quad (1)$$

A 4004 die built on a modern process would be the size of a sawdust particle, and just a few times larger than a red blood cell.

The manufacturing cost of deep-submicron silicon, in volume, is about \$16 / square inch, so each 4004 microprocessor would cost about 0.0017 cents to produce. A few-thousandths of a cent per processor is a low enough price to think about building a few more than strictly needed; in fact, you could go ahead and build one million microprocessors for just seventeen dollars in manufacturing cost.

A jar containing one million sawdust-particle-sized microprocessors might be fun to have on your desk and show to your friends --- well worth the cost just for that purpose --- but it wouldn't be useful for much else.

To make those useless microprocessors into useful programmable matter particles, you would need to add a power supply, memory, communications transceiver, and sensors or actuators to each one. How to accomplish that is the subject of this thesis.

1.4 Scope

It is our hypothesis that the materials and fabrication process steps to construct programmable matter already exist; that the problem of constructing programmable matter is chiefly a problem of integrating already-demonstrated devices into working autonomous microsystems. This thesis is the result of my efforts to design and build working programmable matter particles. I approached this task by sketching out designs, computing key engineering quantities, testing promising designs with physical simulation, and finally building centimeter-scale functional prototypes.

In the course of doing this work, I discovered that a lot of fundamental questions about the hardware design of programmable matter particles had yet to be answered. In particular: How much energy and power are available on a particle? How much is required for applications? How much energy does it take to send a bit a given distance? How much force can a particle exert on its neighbors? How do all of these quantities scale with particle size? The answers to these questions are the subject of the first part of this thesis, called "Physics, Materials, and Devices."

In the second part, "Application Examples," I present what I believe are workable basic designs for a paintable display and programmable shape-change system, and recount some lessons we learned from building and working with centimeter-scale prototypes.

1.5 Background

This thesis concerns hardware implementation of the architectural visions of amorphous computing [12], paintable computing [13] and programmable matter. [14]. The application examples are in displays and modular robotics; for an excellent survey article on the state-of-the-art in modular robotics, see [15].

The computation rate and memory capacity required to achieve reasonable functionality in a programmable matter system is a key design consideration, especially as the node size becomes small. Fully distributed PostScript rendering for a paintable display can be accomplished scaleably and efficiently, as is shown later in this document. Optimal self-reconfiguration planning for modular robots [16] is unlikely to be solvable in polynomial time [17], so workers use simulated annealing [18], genetic algorithms [19], or draw inspiration from biology [20], physics [13], or human cognition [21] to design algorithms to accomplish the task in a suboptimal but bounded-resource manner. Self-reconfiguration control for sliding-cube geometry modular robots (such as those to be considered in detail in the application examples section) can be accomplished using cellular automata. [22]

Modular robot hardware [23], such as Polybot [24] , CONRO [25], M-Tran [26], Molecule, [27], Crystal [28], the I-cube [29], and the molecubes [30] generally have nodes that are greater than 10 cm in smallest dimension and cost more than \$50 per node to produce. This is because they are typically made using off-the shelf, macroscale electronic and mechanical components, and a large number of standardized components are required to produce a functional system.

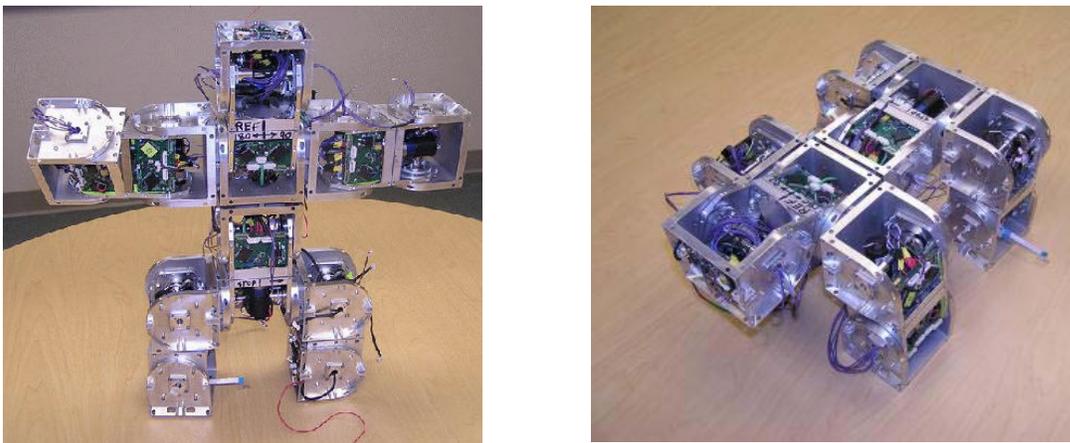


Figure 6: Shen's SuperBot configured as a biped walker (**left**) and as a quadruped walker. (**right**) [31]

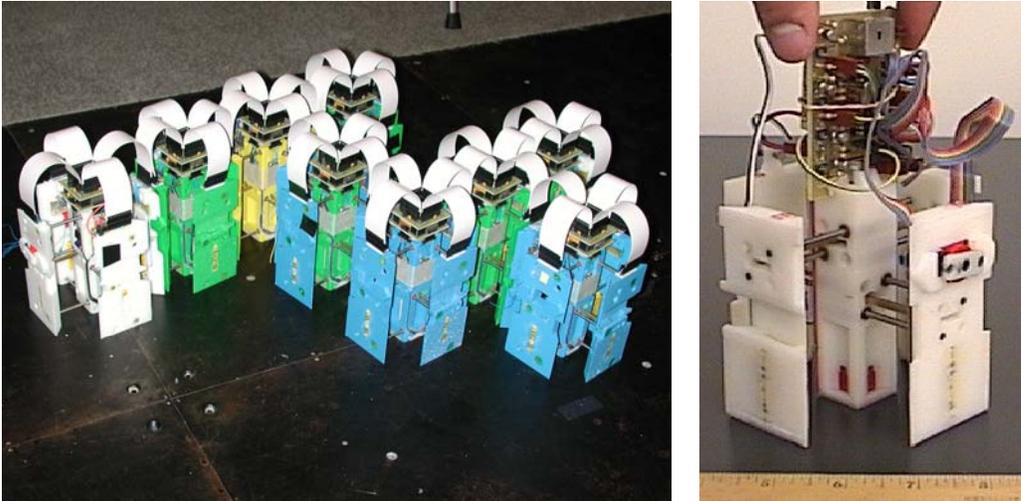


Figure 7: Vona and Rus's Crystalline robot (left) has unit-compressible modules (right) that can change size by a factor of two and latch for reconfiguration and amoeba-like mobility. [28]



Figure 8: Zykov and Lipson's Molecubes swivel about their axis and use magnetic-latching to achieve self-reconfiguration. [30]

To make modular robots a useful raw material for building products, the per-node cost must be substantially reduced. Over the past several years, modular robot design has moved toward systems with few or no moving parts in the nodes. Some of these systems utilize an external fluid bath and external agitation to provide the force and energy to make and break connections [32], [33], [34], controlling node-to-node adhesion to steer the structure toward the desired result. Kirby and his colleagues [35] describe 24 mm diameter cylindrical nodes, capable of translating in a plane by rotating around one another by activating a radially positioned array of electromagnets. Another strategy is employed by the Miche self-disassembling modular robot [36], which starts with all nodes connected, and then releases magnetic latches to disconnect

nodes that are not part of the structure. These systems have lower per-node cost and are more amenable to microfabrication than the previous generation of designs.

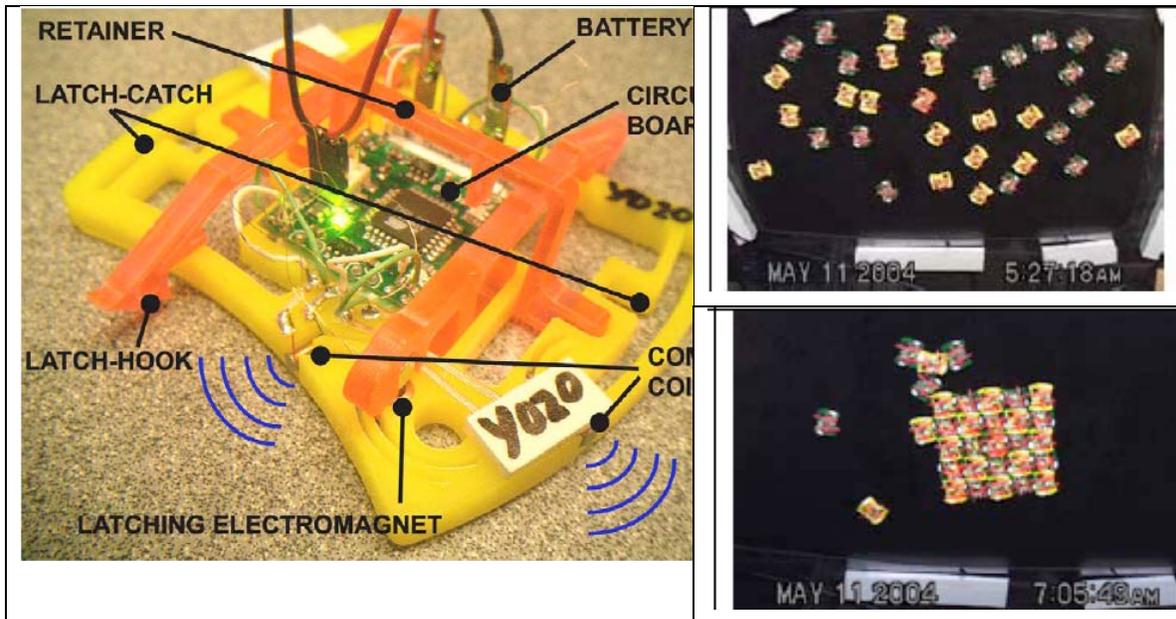


Figure 9: Griffith and Jacobson's latching tiles (**left**) start in a random configuration on an air table (**upper right**) and stochastically reconfigure through random agitation. When a node finds itself in the correct position, it latches to its neighbors. Over 90 minutes, the nodes reconfigure into an ordered configuration. (**lower right**) [33]

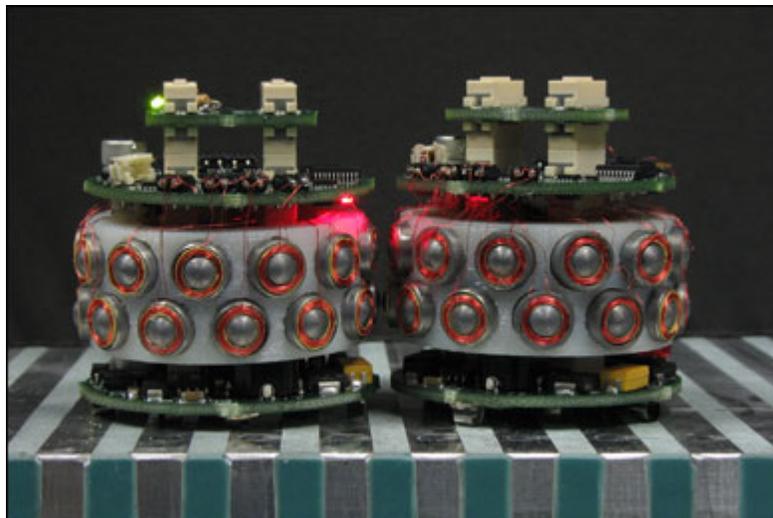


Figure 10: Kirby and Goldstein's catoms rotate relative to one another by activating radial electromagnets. [35]

Autonomous microsystems are integrated circuits with on-board power conversion and communications circuitry. Prior art in this area includes systems that are completely wireless and autonomous, such as the Berkeley Smart Dust motes, [37] and nodes that self-assemble into a regular grid, such as the drivers for the Alien Technologies CMOS nanoblock display [38].

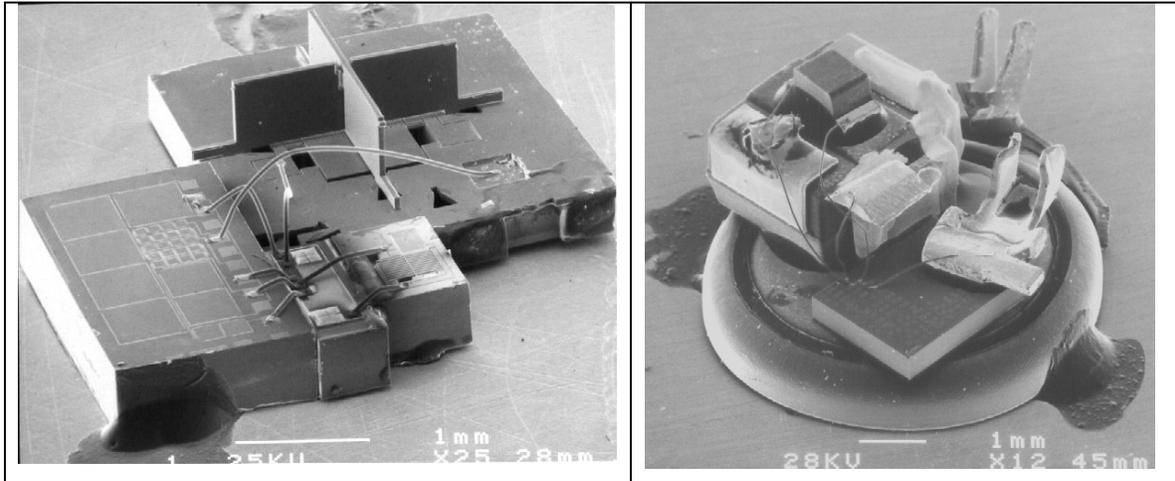


Figure 11: U.C. Berkeley several-millimeter-scale “Smart Dust” sensor network nodes. A 4.8 mm^3 solar-powered node with bidirectional free-space-optical communication, acceleration and ambient light sensing (**left**) and a 138 mm^3 battery powered node with unidirectional communication and ambient light sensing. (**right**) [37]

2 PHYSICS, MATERIALS, AND DEVICES

It is not always advisable to use our everyday intuition to design systems whose size is very large (e.g. rockets, skyscrapers) or very small (e.g. MEMS devices) compared to our everyday experience. Many physical phenomena are scale dependent. Water striders can walk on water, but elephants cannot, because surface tension becomes stronger than gravitation with decreasing scale.

The autonomous electronic and mechanical devices (e.g. laptop computers, digital music players, automobiles) that we live and work with typically have a power source which many orders of magnitude larger in mass and volume than their computational elements. For example, a typical laptop computer has a battery with a mass of several kilograms, but its unpackaged microprocessor and memory have a mass of only a few grams.

To outfit grains of sand with computational power, the power source, communications antenna, heat dissipation surface, and actuators will be many orders of magnitude smaller than we are used to; our hard-won engineering intuition about what is possible and what is not, both in mechanics and electronics, may not apply.

Some of the intuition needed to design a programmable matter particle can be gained by studying the examples in microelectronic [39] or micromechanical [40] system design texts. However, these works do not consider the design of autonomous microsystems, which present their own special scaling issues. In this section, I attempt to fill that gap through literature references and first-principles physical analysis.

To be able to design a millimeter-scale programmable matter particle, we needed to answer the following questions:

- How much power is available from various conceivable sources?
- How much power is required?
- How much heat can be dissipated by a particle?
- What kind of inter-particle communications system should be used?
- How much energy per bit is required for inter-particle communications?
- How will particles be manufactured?
- How much will it cost to manufacture particles?

This section attempts to answer these questions. We hope that it will be a useful resource for anyone beginning the design for a programmable matter particle.

2.1 Heat

Heat dissipation is a critical factor in the design of particles containing magnetic actuators, light emitters, or high-performance computers.

Heat dissipation per volume increases as length scale goes down. This is because heat dissipation is approximately an area effect, and surface area to volume ratio goes up as length scale goes down. From the programmable matter system designer's perspective, this is the good news. It means that breaking a system into many small pieces and spreading them around allows it to run cooler.

However, the amount of heat that a single particle can dissipate certainly does go down as the length scale of the particle goes down. This section shows how much heat is cause for concern at a given length scale, from 10 μ m through 10 cm. Engineering charts of power vs. length scale vs. particle temperature are presented, for continuous and pulse-mode operation. Using these charts, the designer can get an idea of how much power dissipation is reasonable at a given length scale.

These charts are based on a model for natural convection around a sphere. In other words, they do not take the geometry of the particle into account. Heat fins, actuators to blow air around, and other engineered heat transfer solutions can increase the allowable heat dissipation beyond that predicted by the charts in this section. Still, these charts can be used for setting specifications for a first-pass design.

2.1.1 Heat Transfer and Thermal Resistance

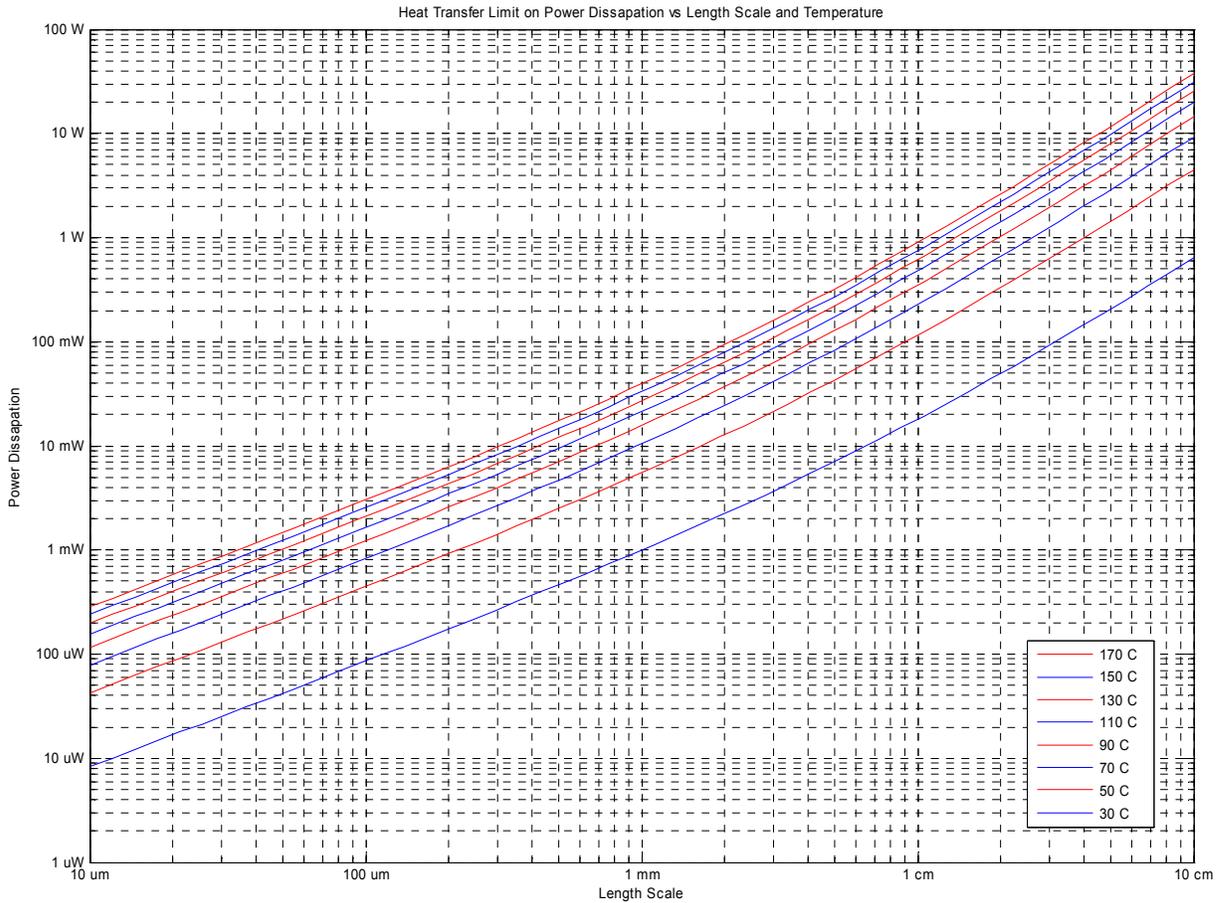


Figure 12: Heat Transfer Limit on Power Dissipation: This chart can be used to determine how much power dissipation is reasonable, given a length scale and maximum operating temperature. These are approximate values; see text for detail. From the plot, one can see that a 1-mm device at 70°C can dissipate about 10mW. (A 25°C free-air temperature is assumed for this plot.)

Figure 12 gives a first approximation for the steady-state operating temperature of a particle given its power dissipation and size. For example, using the table, one can see that a 1-mm device can dissipate up to 10 mW before reaching 70°C, whereas a 100 μ m device can only dissipate about 800 μ W. Figure 13 shows the same data presented in thermal resistance format.

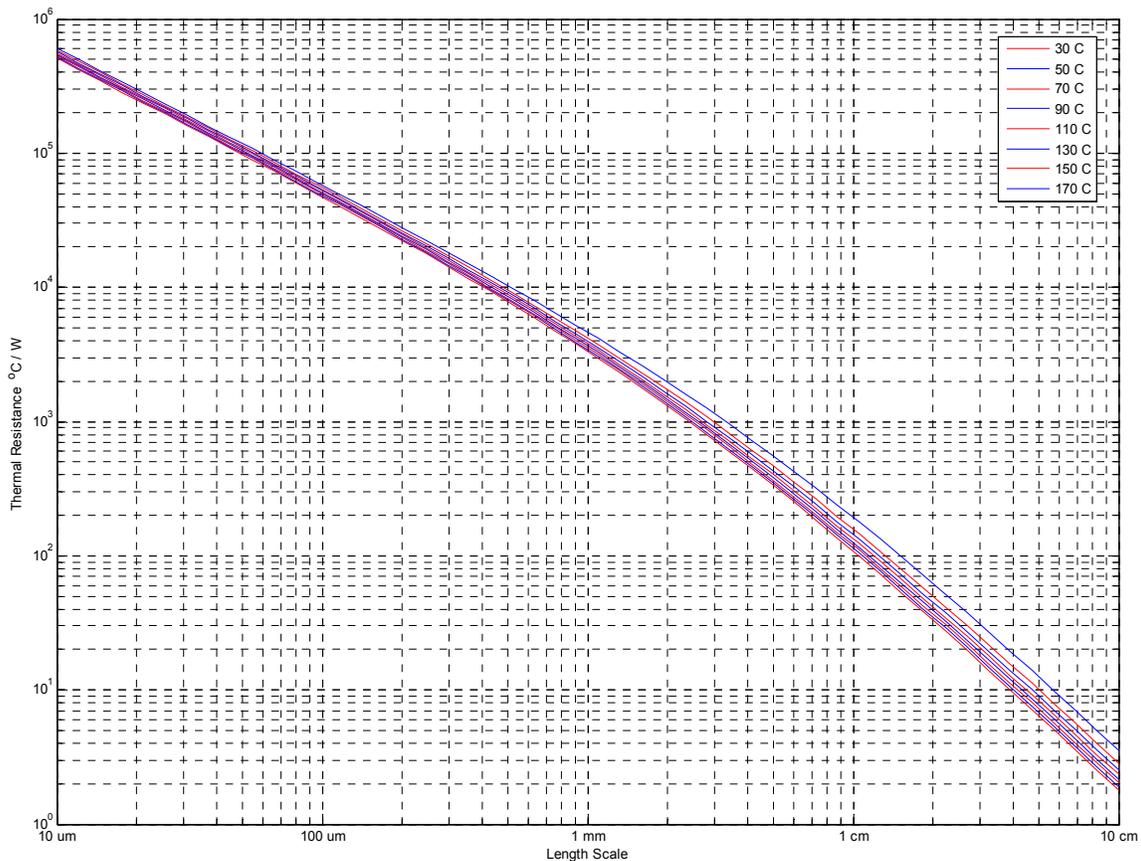


Figure 13: Thermal Resistance: This is the same information as Figure 12, presented in thermal resistance format, which can be more useful for calculation. These are approximate values; see text for detail. From the plot, one can see that a typical 1-mm device has a thermal resistance of about $4000\ ^{\circ}\text{C}/\text{W}$.

Figure 12 and Figure 13 show the heat transfer by natural convection and radiation from a sphere of a given diameter, at a given temperature. It is possible to achieve more heat transfer than shown by using a higher surface area shape than a sphere, increasing the surface area with fins or by adding actuators to move air. It is also possible to get lower heat transfer, by using a low surface area shape, or by coating the particle with a material that is not sufficiently conductive of heat.

T (K)	Sphere surface temperature
D (m)	Sphere diameter
$T_{amb} = 298$ K	Free-air temperature
$T_{film} = \frac{T + T_{amb}}{2}$	Film temperature
$\beta = \frac{1}{T_{film}}$	Coefficient of volumetric expansion (for an ideal gas)
$g = 9.81$ m/sec ²	Acceleration of Gravity
$\nu = fcn(T_{film}) \approx 2 \times 10^{-5}$ m ² /sec	Kinematic Viscosity of Air
$Gr_D = \frac{(bouyancy)(inertia)}{(viscosity)^2} = \frac{\beta(T - T_{amb})gD^3}{\nu^2}$	Grashof Number
$Pr = 0.71$	Prandtl Number of Air
$Ra_D = \frac{(bouyancy)}{(viscosity)} = Gr_D \cdot Pr$	Raleigh Number
$Nu_D = 2 + \frac{0.589 \cdot Ra_D^{1/4}}{[1 + (0.469/Pr)^{9/16}]^{4/9}}; Ra \leq 10^{11}; Pr > 0.5$	Nusselt Number Correlation, "Natural Convection on a Sphere"
$k = fcn(T_{film}) \approx 0.26$ W/m·K	Thermal Conductivity of Air
$h_c = \frac{Nu_D k}{D}$	Convection Heat Transfer Coefficient
$A = \pi \cdot D^2$	Surface area of a sphere
$\dot{Q}_c = h_c A(T - T_{amb})$	Convection Heat Transfer Rate
$\varepsilon = 0.80$	Emissivity of Sphere (Value for black anodized aluminum; typical value for non-reflective engineering materials)
$\sigma = 5.670 \times 10^{-8}$ JK ⁻⁴ m ⁻² s ⁻¹	Stefan-Boltzmann Constant
$\dot{Q}_r = \varepsilon A \sigma (T^4 - T_{amb}^4)$	Radiation Heat Transfer Rate
$\dot{Q} = \dot{Q}_c + \dot{Q}_r$	Total Heat Transfer Rate
$R_{th} = (T - T_{amb}) / \dot{Q}$	Total Thermal Resistance

Table 1: Calculating Heat Transfer and Thermal Resistance

Formulas and constants used in Table 1 are taken from [41].

2.1.2 Thermal Capacitance and Thermal Time Constant

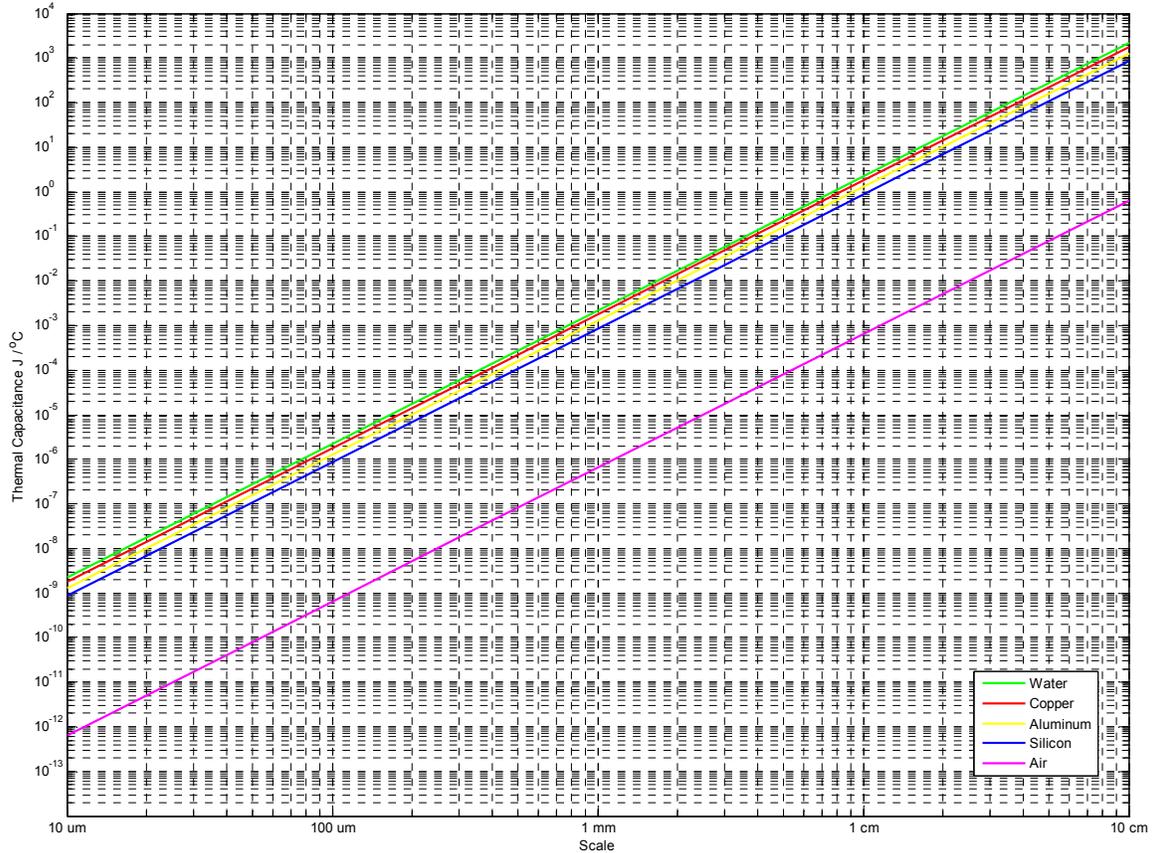


Figure 14: Thermal Capacitance: This chart may be used to estimate the thermal capacitance of a programmable matter particle. It shows the thermal capacitance of a sphere with a given diameter, for a variety of materials. From the plot, we see that a 1-mm silicon device has a thermal capacitance of about $2 \text{ mJ} / ^\circ C$.

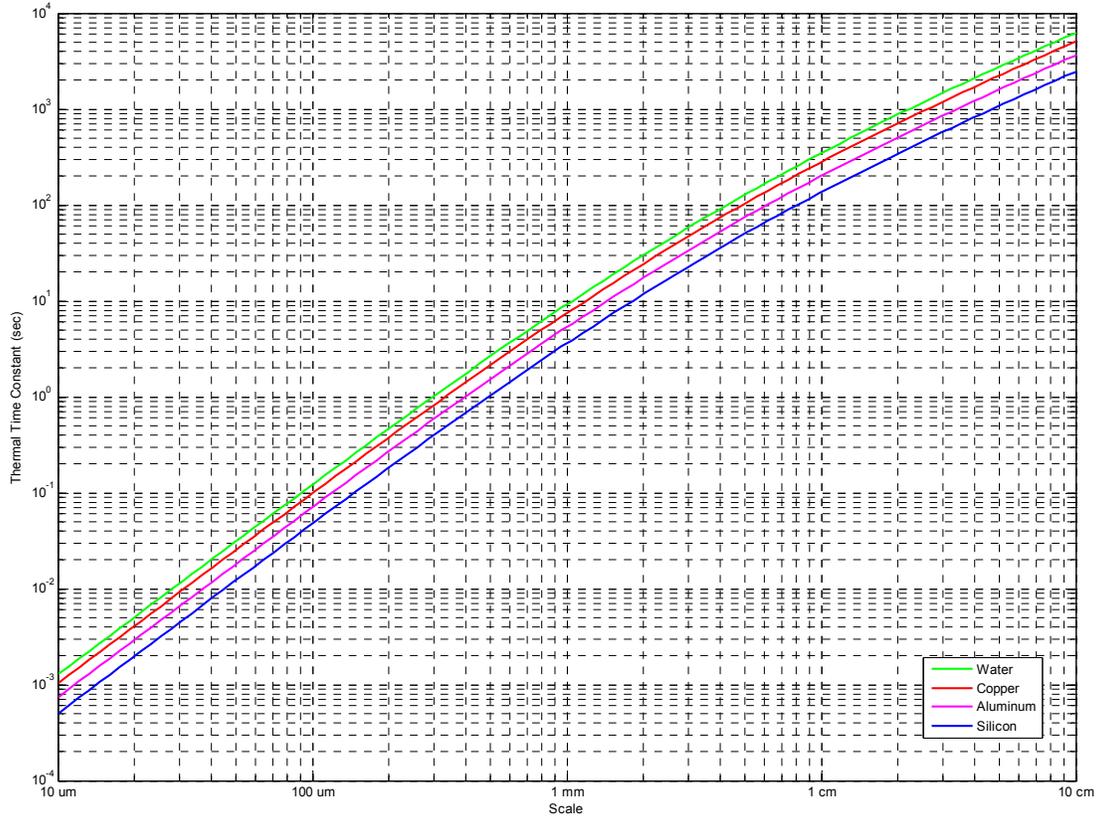


Figure 15: Thermal Time Constant: This chart may be used to estimate the thermal time constant of a programmable matter particle. This chart is subject to the same approximations as Figures 1-3: that the system is a sphere of a given diameter. From the chart, we can see that a 1-mm system has a thermal time constant from 3-8 seconds.

R_{th}	Total Thermal Resistance
$V = \frac{4}{3} \pi \cdot D^3$	Volume of a sphere
$\rho_{Cu} = 8920 \text{ kg/m}^3$ $\rho_{Si} = 2330 \text{ kg/m}^3$ $\rho_{H_2O} = 1000 \text{ kg/m}^3$ $\rho_{Al} = 2700 \text{ kg/m}^3$	Density of Materials
$H_{Cu} = 380 \text{ J/kg-K}$ $H_{Si} = 705 \text{ J/kg-K}$ $H_{H_2O} = 4186 \text{ J/kg-K}$ $H_{Al} = 900 \text{ J/kg-K}$	Specific Heat Capacity of Materials
$C_{th} = \rho V H$	Thermal Capacitance
$\tau_{th} = R_{th} C_{th}$	Thermal Time Constant

Table 2: Calculating Thermal Capacitance and Time Constant

2.1.3 Pulsed Power Safe Operating Area Curves

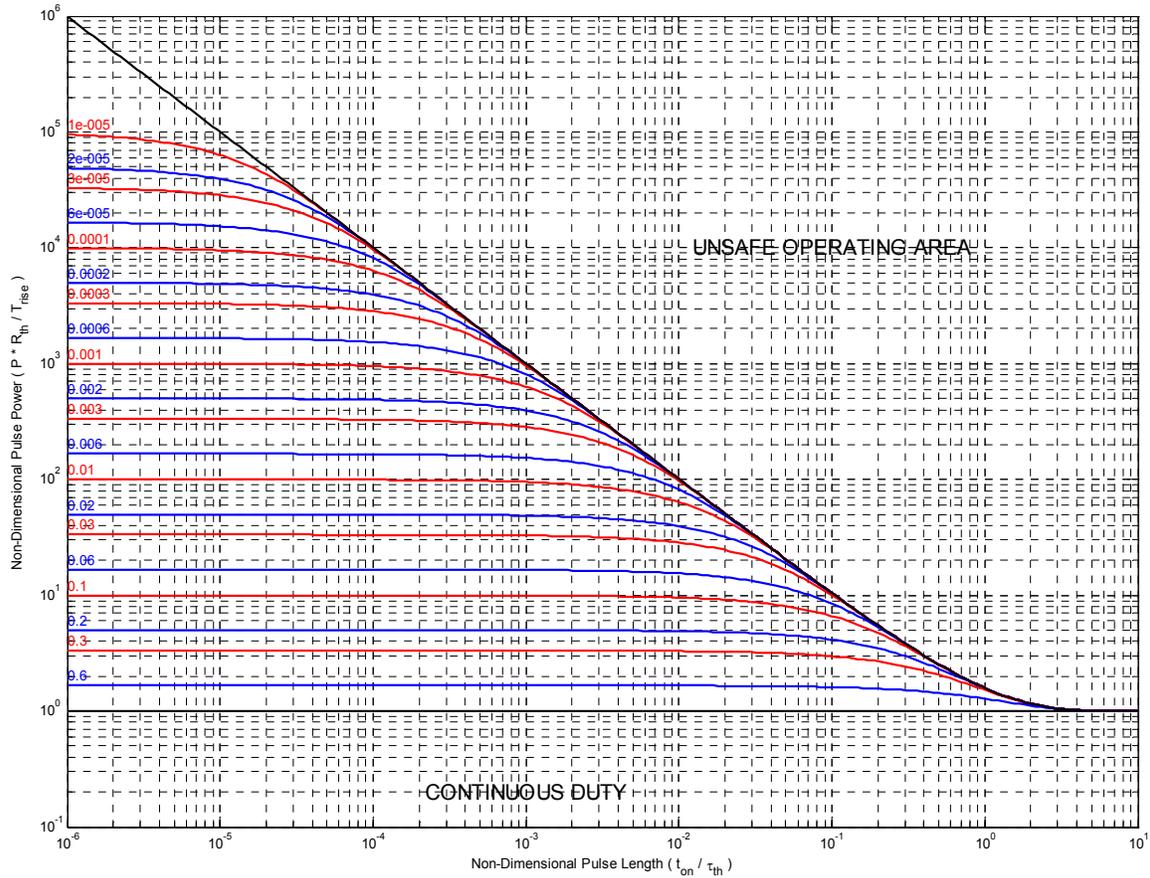


Figure 16: Thermal SOA for Pulsed Operation (Non-Dimensional): This chart shows the generalized relationship between pulse length, thermal time constant, pulse power, thermal resistance, maximum temperature rise, and maximum duty cycle. In the lower region, continuous operation is allowed; in the upper region, a single pulse will cause the system to overheat; in the middle region, operation is allowed so long as the duty cycle is equal to or lower than that indicated.

Figure 14 and Figure 15 provide a way to estimate the thermal capacitance and thermal time constant of a programmable matter particle. The thermal capacitance is the amount of thermal energy that a system can store per degree increase in temperature. The thermal time constant is a measure of the rate at which a system cools to ambient temperature. These are important figures in the design of a system that dissipates power in a pulse mode.

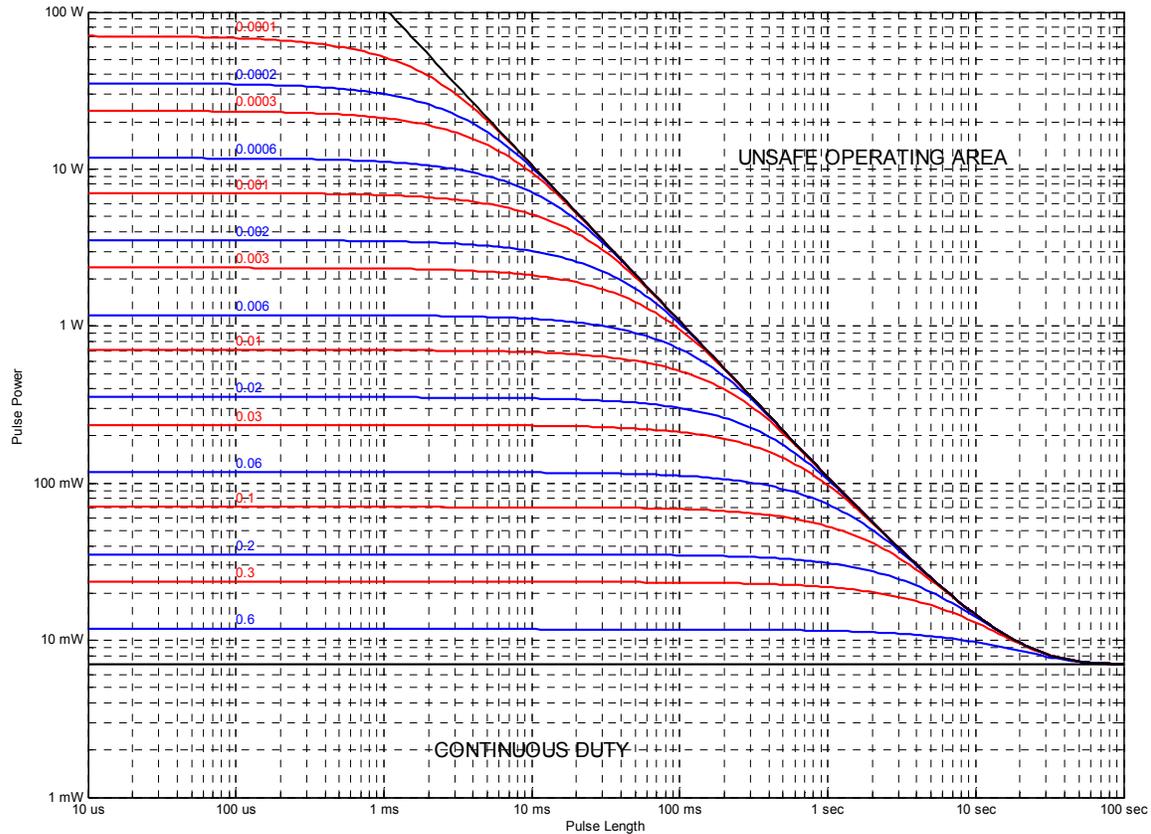


Figure 17: Thermal SOA Plot for Pulsed Operation, 60°C rise, 1-mm: This chart shows the maximum duty cycle for pulsed power dissipation from a 1-mm programmable matter particle. Reading the chart, one can see that a 1-mm particle operating with 1-sec 30mW pulses (e.g. from an RF transmitter) must cool down for about four seconds between pulses (i.e. a duty cycle of 20%) to avoid overheating. This chart is drawn for a copper sphere in free air experiencing uniform volumetric heating.

T	Particle Temperature
T _a	Ambient Temperature
T	Time
P(t)	Instantaneous System Power Dissipation
$\frac{dT}{dt} = C_{th} [P(t) - (T - T_a)R_{th}]$	Differential Equation describing particle thermal model
P _{pulse}	Pulse Power
t _{on}	Pulse Length
DC	Duty Cycle
R _{th}	Particle-to-Ambient thermal resistance
T _{th}	Particle thermal time constant
$\frac{P_{pulse} R_{th}}{\Delta T_{rise}} = \frac{1 - e^{-(t_{on}/DC)/\tau_{th}}}{1 - e^{-t_{on}/\tau_{th}}}$	Solution to differential equation describing system model: relationship shown in Figures 5 and 6

Table 3: Calculating Thermal Safe Operating Areas

2.2 Power

Physically, programmable matter consists of autonomous microsystems, called particles, which are suspended in a fluid or gel, called the binder. The particles require power to function. The ways that one might provide power to the particles can be divided into three categories:

1. Particle-Stored Energy: The energy to operate a particle over its lifetime is stored inside the particle. The energy source is never replaced: once depleted, the system is discarded.
2. Binder-Stored Energy: The energy to operate the particles is stored in the binder. Each particle contains an energy conversion device, but the binder contains the energy source.
3. External Power: The power required to operate the particles is continuously supplied to the system, from an external, usually macroscopic source.

A very large fraction of systems made or sold today use external power, since they can almost always be plugged in, refueled, or have their batteries replaced. For example, automobiles, laptop computers, and homes all fall into this category. Examples of systems that use something akin to particle-stored power include rockets to lift payloads into outer space, glow-sticks, and cheap disposable flashlights with non-replaceable batteries.

If operation in a random environment is desired, the most straightforward solution is to store energy inside each particle. However, the economics of wafer fabrication dictate that cost is proportional to area, regardless of the complexity of that area. Because power sources tend to be large compared to other system components, the particle-stored energy solution places severe restrictions on at least one of the following: device power, device lifetime, and device cost.

Binder-stored energy decouples the volume of the particles from the volume of the energy source, by storing the energy outside of the particles, but still inside of the system. This solution allows relatively high-power, long-lifetime, low-cost systems, compared to the particle-stored power solution. Some materials development will be required to implement binder-stored energy; we know of no prototypes or demonstrations of this concept to date. However, it appears possible in principle.

When continuous operation is required, external power is the only feasible solution. There are some external power solutions that may work in a random environment, (e.g. photovoltaic cells) however, most require a structured environment.

2.2.1 Particle-Stored Energy

2.2.1.1 Electrochemical Cells

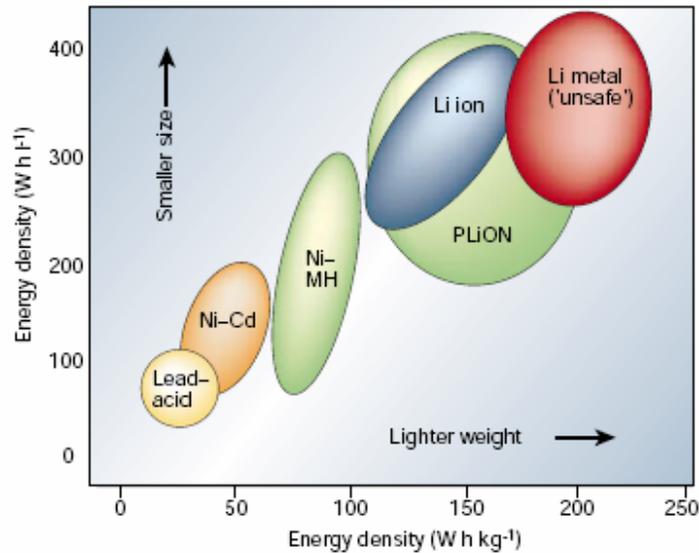


Figure 18: Energy Density of some Battery Chemistries (from [42])

Zinc-Air batteries (which use oxygen from the air as one of the reactants) have the highest gravimetric and volumetric energy density of any battery chemistry, $6.0 \times 10^9 \text{ J/m}^3$ (1500 Wh/L) [43] Lithium-Ion batteries have the highest volumetric energy density of any stable rechargeable battery chemistry, $1.6 \times 10^9 \text{ J/m}^3$ (400 Wh/L) [42]

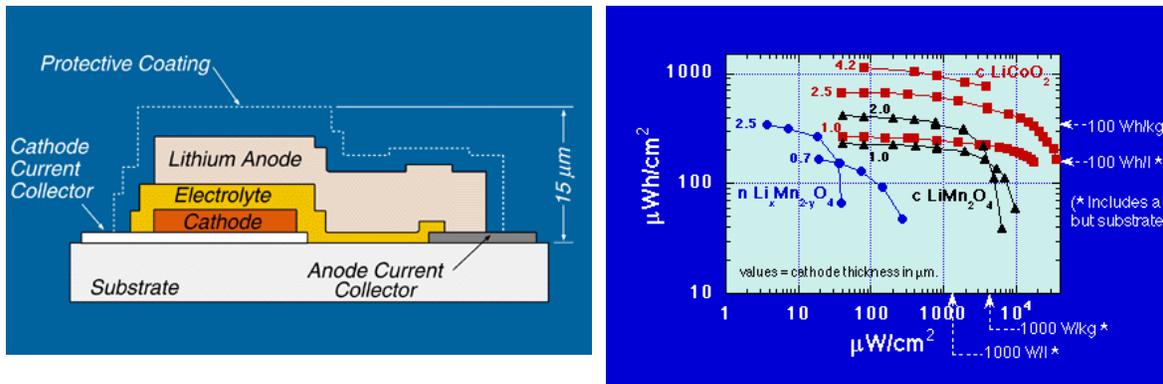


Figure 19: Microfabricated Rechargeable Batteries: A schematic drawing of the microfabricated battery described in [44] (left), and performance curves for thin-film rechargeable battery chemistries. (right) Both figures are from Oak Ridge National Laboratory.

By volume scaling, we can estimate that a 1 mm^3 zinc-air battery could about store **6.0 J**. This is enough energy to run a $2 \text{ }\mu\text{W}$ sensing application for about 34 days, or a 10 mW heat-limited application for about 10 minutes.

Energy densities of microfabricated batteries from the literature support these estimates. In [44], workers at Oak Ridge National Laboratory report fabrication of a 1 cm^2 thin-film rechargeable Lithium battery, using photolithographic techniques, with an energy density of $2.1 \times 10^9 \text{ J/m}^3$.

2.2.1.2 Microengines

Gasoline has an energy density of $2.9 \times 10^{10} \text{ J/m}^3$, making it one of the densest non-nuclear energy storage options available. In macroscopic devices, the energy in gasoline is converted to electricity by combustion.

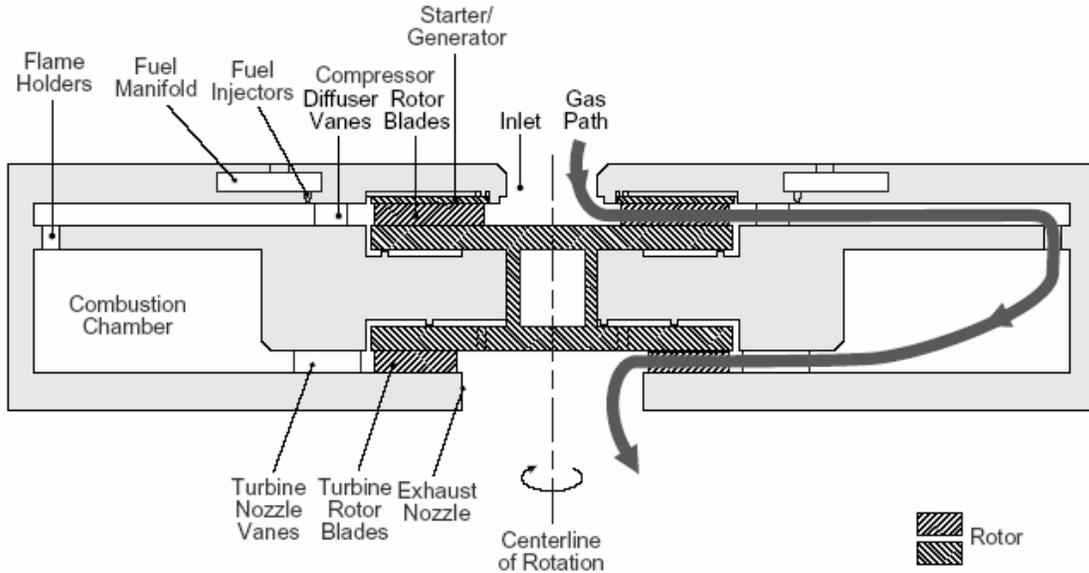


Figure 20: Micro Gas Turbine Generator Cross-Section (from [45])

The fabrication of millimeter-scale combustion engines is a topic of active research. To cite one example, the MIT Microengine Project [45] has built several prototype silicon gas turbine generators, which measure 4 mm on a side. The authors of [45] estimate that their engine will produce $10\text{-}20 \text{ W}$ of electrical power while consuming 10 g/hr of hydrogen. By our calculation, this corresponds to a projected efficiency of $2.5\% - 5\%$.

Using the above numbers, we can make a rough estimate of the effective electrical energy density of gasoline, $1.5 \times 10^9 \text{ J/m}^3$. This is about the same as the current energy density of microfabricated lithium-ion batteries. However, lithium-ion batteries are a very mature

technology, while microengines are very new; it is likely that the efficiency of microengines will increase, perhaps dramatically, with further development.

A system powered by a microengine will be able to operate at much higher power levels than a battery powered system, due to forced convection of the exhaust. However, the presence of an exhaust stream from every particle of programmable matter might become irritating or dangerous to the operator, depending on the fuel selected.

We can estimate that 1 mm³ of gasoline plus a microengine can store and convert **1.5 J**. This is enough energy to run a 2 μW sensing application for about 8 days, or a 10 mW application for about 2.5 minutes.

2.2.1.3 Radioisotopes

Radioisotopes can have tremendous volumetric energy density. ¹⁸⁰Ta, which has a half-life of eight hours, has an energy density of **1.0 x 10¹⁵ J/m³**, about 30,000 times the energy density of gasoline. A 1 mm³ sample of ¹⁸⁰Ta would continuously release 34 W. ¹⁷⁸Hf, which has a half-life of 31 years, has an energy density of **1.0 x 10¹⁶ J/m³**, about 300,000 times that of gasoline. A 1 mm³ sample would release 160 mW. [46]

However, two major factors currently restrict the usefulness of radioisotope power sources in programmable matter: heat production and toxicity to humans.

Radioisotope power may be appropriate for systems that:

- A. Require a very small amount of power
- B. Must operate continuously for a very long time (e.g. 10-100 years)
- C. Cannot be accidentally ingested or inhaled
- D. Can be thoroughly gathered up and properly disposed of at end-of-life

2.2.1.3.1 Heat Production

To understand the heat production issue, suppose that a programmable matter particle were powered by a 1 mm³ sample of ¹⁷⁸Hf, and that all of the energy released by that sample was eventually converted to heat inside the particle. A 1 mm³ particle has a 10 mW heat dissipation limit by natural convection, (see Figure 12) so the 160 mW emitted by the sample of ¹⁷⁸Hf would cause the particle to melt before leaving the factory.

To avoid this problem, a small enough volume of ¹⁷⁸Hf would need to be designed into the particle so that at most 10 mW of power was being released. However, microscale radioisotope energy conversion, though betavoltaic devices, tends to have a very low efficiency, about 1%. [47] With 1% conversion efficiency on the 10mW released, only 100μW of electrical energy would be available for use by the particle.

Unless the efficiency of microscale radioisotope energy conversion can be increased very substantially, or unless significant active cooling is used, radioisotopes cannot provide high power densities to programmable matter particles, even before considering toxicity concerns. However, radioisotope sources are capable of very high energy density, for very long life, very low power applications.

To continue our comparison from the previous section, *based on heat considerations alone*, a radioisotope source could power a 2 μ W sensing application for many years, but could not power a 10 mW application at all.

2.2.1.3.2 Toxicity

The main radioactive isotopes under active investigation for microscale power, ^{63}Ni and ^3H , are beta-particle emitters. Beta particles are electrons. The beta particles from these two isotopes have a low enough energy that they do not penetrate the outer layer of dead skin, or travel through more than a few inches of air. [48, 49]

However, if ingested or inhaled, beta particle emitters can still be highly toxic to humans, and can cause genetic damage, cancer, radiation sickness, and death. In many applications, including most of those listed in the introduction, there may be the potential for programmable matter to be accidentally ingested or inhaled.

The primary vector for radiation dose from ^{63}Ni is uptake into the bone, by inhalation of vapors or direct ingestion. [48] Ingestion of 20 μ Ci/year of ^{63}Ni causes a radiation dose to the bone of 0.01 REM/year.

The U.S. Nuclear Regulatory Commission sets the dose limit to any part of the body to 0.1 REM/year for the general public, and 10 REM over five years for specially licensed *nuclear energy workers*. The occupational dose limit is liberal; if routinely exposed to the maximum dose over the course of a career, one would expect 25% of nuclear energy workers to die of cancer, versus 20% of the general population. [50, 51]

Scaling the ^{63}Ni dosimetry information, we can see that the NRC exposure limit for a member of the general public allows the ingestion or inhalation of a maximum of 200 μ Ci/year of ^{63}Ni , assuming that this was this person's only exposure to man-made radiation.

A prototype betavoltaic cell has been constructed using a ^{63}Ni source. [52] Extrapolating from experiments with lower-activity samples, the authors of [52] report that it should be possible to produce 150 nW of continuous electrical power from a 100 mCi source of ^{63}Ni . This is a very radioactive source; if a person were to accidentally ingest it, they could receive a radiation dose *500 times larger* than the annual dose limit for the general public, and *25 times larger* than the annual occupational dose limit.

Still, microscale betavoltaic batteries are a very new technology, and it is possible that they will improve in efficiency, perhaps by orders of magnitude, with further development. Also, it may be possible to encapsulate the radioactive material inside a very durable protective coating, so that a person would not absorb the radioactive material into their body, even if they did ingest programmable matter.

2.2.2 Binder-Stored Energy

Binder-stored energy solutions store reactants or fuel in the binder, around the particles, rather than inside the particles. This allows relatively high-energy, long-lifetime, low-cost systems, compared to systems with particle-stored power.

Some materials development will be required to implement binder-stored energy in programmable matter; we know of no prototypes or demonstrations of this concept to date.

2.2.2.1 Inside-Out Zinc Air Battery

An inside-out zinc-air battery, for a spray-on display application, is shown in Figure 21. With the inside-out zinc-air battery concept, the battery reactants are stored in the binder, rather than in the particles. Because of comparatively large volume of the binder, this results in longer battery life.

Consider a paintable display with a paint thickness of 1 mm, and with cubic 110 μm particles, each drawing 50 μW , with one particle for every 0.25 mm^2 area. (These are the specifications for the 640 x 480, 17" paintable display described in the application examples.)

If particle-stored power were used, then a maximum of 0.0013 mm^3 would be available for energy storage inside each particle. This would be enough volume to store 8 mJ of energy using a zinc-air battery, so the battery life of this display would be limited to 2.5 minutes.

With the inside-out zinc air battery, an example of binder-stored power, about 0.25 mm^3 per particle is available for the storage of energy, which is enough volume to store up to 1.5 J per particle. In this case, the battery life of the display could be up to 8 hours.

The zinc and zinc oxide must travel through the polymer matrix in the binder by diffusion. We can get a rough idea of the particle size required to achieve suitable power density by evaluating the diffusion time constant for a zinc particle. The diffusion time constant should be about equal to the desired service lifetime of the battery.

The diffusion time constant for a zinc particle with diffusion coefficient D over a length scale d is given by

$$\tau_D = \frac{d^2}{\pi^2 D} \quad (2)$$

The Stokes-Einstein relation gives the diffusion coefficient of a particle in terms of the Boltzmann constant k , absolute temperature T , zinc particle diameter σ , and the dynamic viscosity of the liquid phase η .

$$D = \frac{kT}{6\pi\eta\sigma} \quad (3)$$

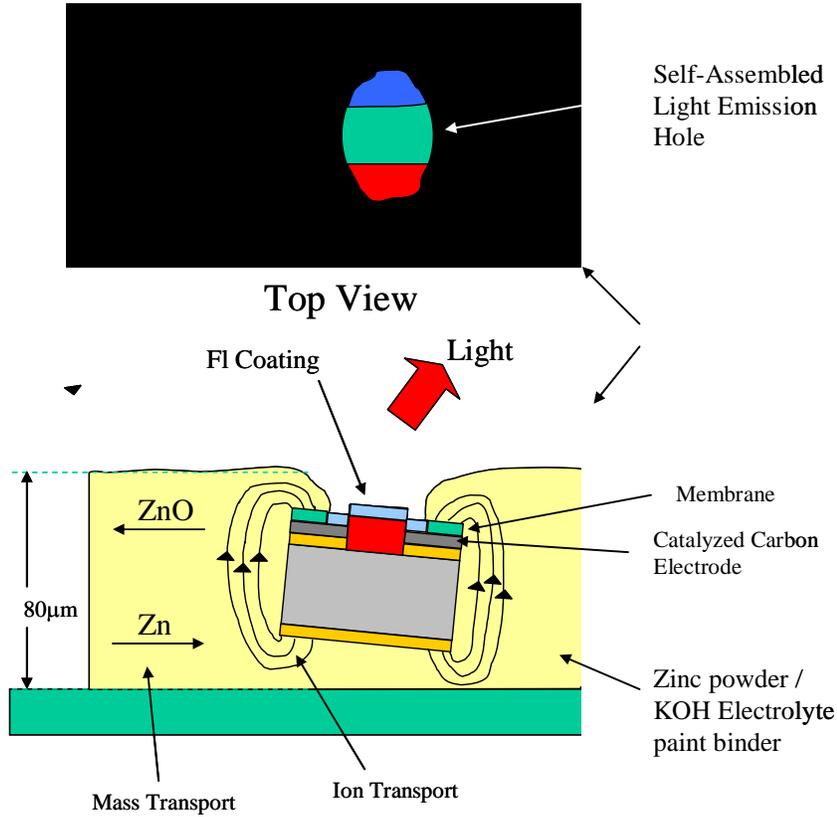


Figure 21: Inside-Out Zinc Air Battery Concept for a Paintable Display: Paint particles consisting of a CMOS chip, LED, and battery electrodes are painted onto a surface in a Zinc powder / KOH electrolyte binder. The paint binder also contains polymers which harden into a porous matrix. Diffusion brings zinc from the binder to the particle, and takes the reaction product, Zinc Oxide, from the electrodes back into the binder.

Combining equations (2) and (3), we can write the required zinc particle size for a given diffusion rate.

$$\sigma = \frac{\pi}{6} kT \frac{\tau}{\eta d^2} \quad (4)$$

Taking $T = 300 \text{ K}$, $\tau_D = 8 \text{ hours}$, $\eta = 8.7 \times 10^{-4} \text{ kg/m-sec}$, the viscosity of water, $d = 1\text{mm}$, and $k = 1.38 \times 10^{-23} \text{ m}^2\text{-kg/s}^2\text{-K}$, the maximum zinc particle radius σ is 72 nm. It is possible to fabricate zinc nano-flakes as small as 3-5 nm in diameter by dry roller vibration milling [53], so it is possible to fabricate zinc particles small enough for the inside-out zinc-air battery discussed here to function by diffusion.

2.2.2.2 Binder-Stored Fuel

In the combustion engine with binder-stored fuel concept, each particle contains a combustion engine which is powered by fuel drawn in by capillary action or diffusion from pores in the binder, and by oxygen from the air. Similar energy densities to the inside-out zinc-air battery are possible, with potentially larger power densities.

2.2.3 External Power

In this section we consider technologies for supplying power to programmable matter from an external source; this power might come from the commercial power grid or from a vehicle's electrical system.

In §2.2.1-§2.2.2, we have shown that it is possible to run programmable matter systems, even relatively power hungry ones, on battery power for several hours. This is on par with the battery life of ordinary macroscopic systems, like laptop computers and cellular telephones. When the batteries in these devices run out, they can be plugged in and recharged. This section is about how to “plug in” a paintable system, either to recharge its batteries, or to allow for continuous operation.

Acceptable methods for supplying power to a paintable system cannot require hand-manipulation of individual particles. Having personnel attach tiny connectors to each of the millions of particles that make up a paintable system would be absurd. Also, power transfer through conductive planes is not ideal, because a single shorted particle could short out the entire system, and because the particles would need to be permanently attached to the planes. What is needed is a hands-off, wireless, batch process for supplying power the sparsely distributed, randomly oriented paint particles.

There are a wide range of possible options; for a general survey of power harvesting techniques for mobile electronics, see [54]. Here we focus on two options that we feel are the most feasible for programmable matter: photovoltaic cells and reactive power transfer.

2.2.3.1 Photovoltaic Cells

Particles powered by photovoltaic cells can operate indefinitely in a random environment, so long as it is not too dark. This is a claim that cannot be made regarding any of the other power sources discussed in this document, with the possible exception of the radioisotope sources.

Since high-efficiency photovoltaic cells are already made using wafer-fabrication technology, little technology development would be required to use them as an energy source for the particles.

The full-sun outdoor solar irradiance is 100 mW/cm^2 . Indoor irradiance is typically less than 1 mW/cm^2 . [55] The best single crystal cells, fabricated using a GaInP / GaAs process, have 30.3% efficiency. Single-crystal silicon cells have an efficiency of 24.7%. [56]

Applying these figures to a 1 mm^2 cell, photovoltaic cells can deliver **300 μW** outdoors, when the sun is shining, and **3.0 μW** indoors, when the lights are on.

Shining a lamp on sparsely-distributed programmable matter particles equipped with photovoltaic cells, in an attempt to transfer power, results in poor efficiency. The most efficient lamps available are sodium-vapor lamps, which have an efficiency of about 34%. We can multiply this by the efficiency of the GaInP / GaAs solar cells to get a direct transfer efficiency of 10%. But then, we need to multiply this efficiency by the area-fill-factor of the particles on a surface to get the overall efficiency, since light that does not hit a particle is wasted. The paintable display application discussed here has an area fill factor of 5%, which results in an overall power transfer efficiency of 0.5%. For an application with higher fill factor, this approach might be acceptable.

2.2.3.2 Reactive Power Transfer

High efficiency wireless power transfer to sparsely distributed, randomly oriented programmable matter on a surface can be realized by placing the system inside the inductor or capacitor of a resonant LC tank circuit. This approach is highly efficient because energy that is not absorbed by a particle during one AC cycle is not wasted; most of it is recaptured and used again on the next cycle.

To test this concept, we constructed a prototype paintable display; using 0603 LED's as paint particles. An 0603 LED is about 1 mm long. The prototype system supplied about 120 μW to each of the LED's.

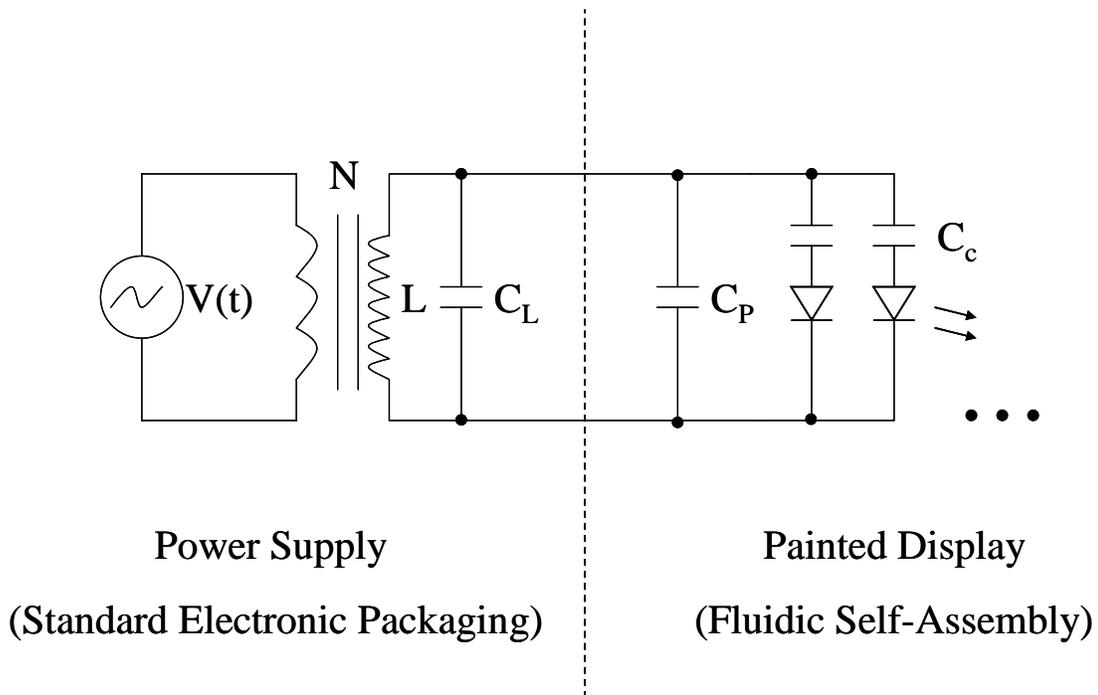


Figure 22: Electrical Model for Reactively Powered Paint: An AC voltage source energizes a transformer-coupled parallel LC tank circuit. The paint nodes (shown as LED's) are placed inside the capacitor, which results in their being capacitively coupled to the parallel LC network.

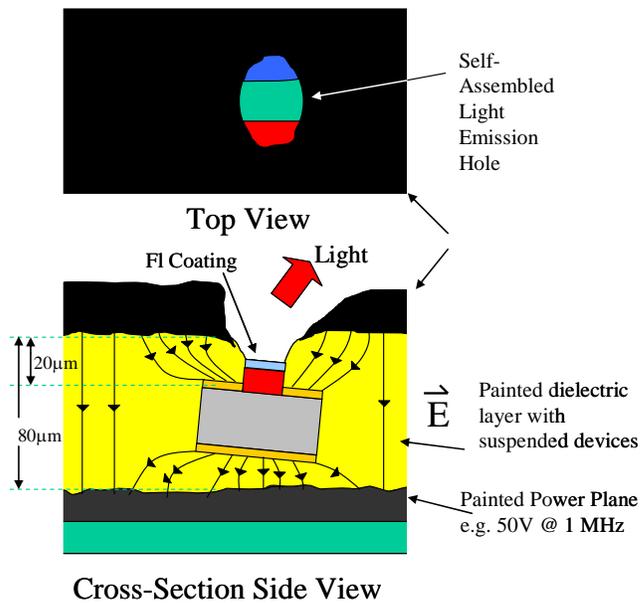


Figure 23: Reactive Power Transfer Concept for a Paintable Display

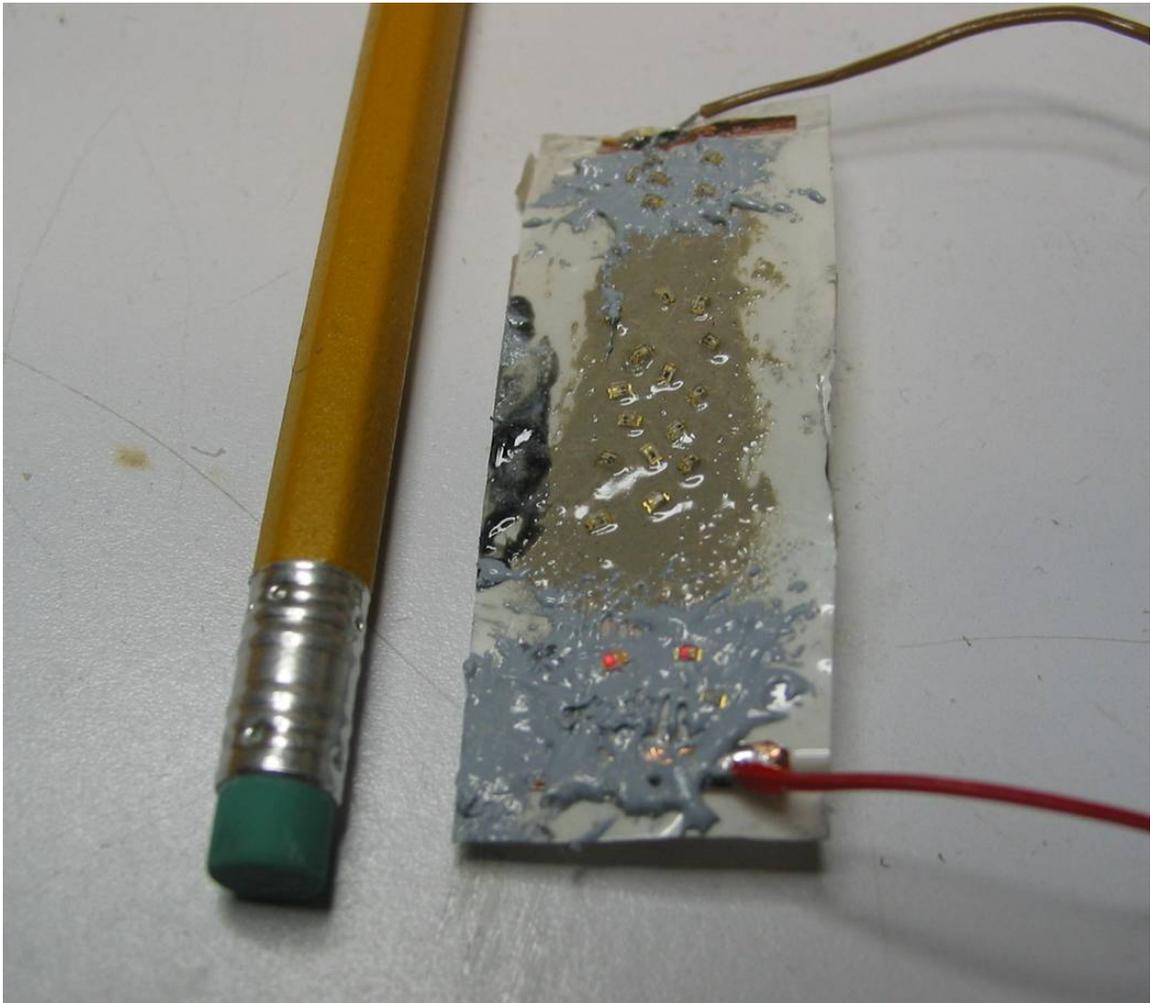


Figure 24: Reactive Power Transfer Demonstration: *We fabricated a single-layer flex circuit with interdigitated electrodes. We coated the circuit with a dielectric paste, made by loading a two-part urethane epoxy to the thickening point with nanophase barium titanate. The paste has a relative dielectric constant of about 9.0. We placed 0603 LED's on top of the paste in random positions and orientations. (The brown paste in the center is from an earlier, non-functional formulation.) We excited the electrodes with 400 VRMS at 500 kHz; two illuminated LED's are clearly visible in color versions of this document.*

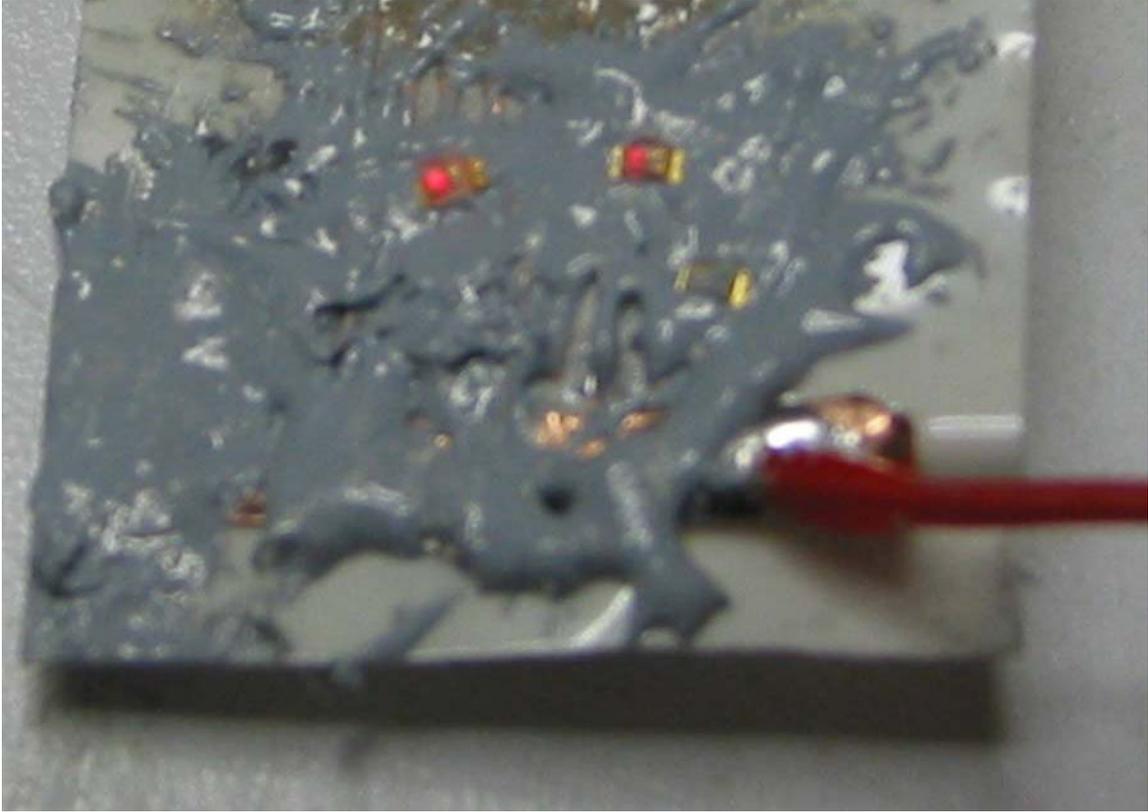


Figure 25: Reactive Power Transfer Demonstration: Close Up

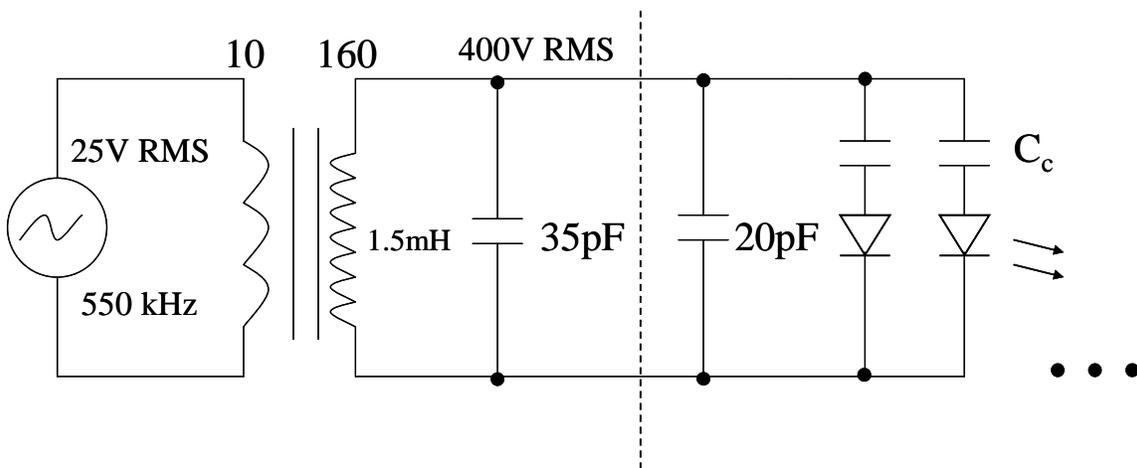


Figure 26: Electrical Details of Reactive Power Transfer Demonstration. A 550 kHz sine wave, produced by a function generator (HP 33120A) is amplified to 25 VRMS by a power op-amp (Apex PA09) and applied to the primary of the transformer. The secondary of the transformer, at 400 VRMS, is connected to a 35 pF capacitor, used for tuning, and to the flex circuit containing the LED's. We measured the capacitance of the flex circuit at 20 pF. From geometry, we estimate the coupling capacitance between the LED pads and flex circuit at 62 fF, which would supply 120 μ W to each LED. This value is commensurate with the level of illumination observed.

It is also possible to place the paint particles inside the inductor of an LC tank circuit. In this case, it would be desirable for the paint binder to be made from a ferromagnetic material. Transparent high permeability materials and polymer composites of these materials exist; see [57]. Also, it would be desirable to place even more strongly ferromagnetic materials inside each particle; for an example of ferromagnetic materials integrated with a wafer-fabrication process, see [58].

2.2.3.3 Power and Efficiency Calculations

From Figure 22, the electrical model for capacitive power transfer is a load in series with a small capacitance for each particle, plus a large parallel shunt capacitance. In this section, we establish approximate mathematical expressions for the maximum achievable power transfer and efficiency to each particle.

The power into a single particle is given by

$$P_p = \frac{V_p^2}{R_p} \quad (5)$$

where P_p is the power transfer to the particle, V_p is the voltage across the particle's terminals, and R_p is the resistance between the particle's terminals. Maximum power transfer from a source to a load occurs when the source impedance equals the load impedance; in this case

$$R_p = \frac{1}{2\pi f C_s} \quad (6)$$

where f is the operating frequency and C_s is the series capacitance to the particle. We can combine equations (5) and (6) to get the maximum power transfer to a particle.

$$(P_p)_{\max} = 2\pi V_p^2 f C_s \quad (7)$$

We can estimate the series capacitance as

$$C_s = \epsilon_o \kappa \frac{A_p}{(L_b - L_p)} \quad (8)$$

where ϵ_o is the permittivity of free space, κ is the relative dielectric constant of the paint binder, A_p is the area of each particle electrode, L_b is the thickness of the paint layer, and L_p is the thickness of the particle.

Combining Equations (7) and (8), we get an expression for the maximum power into a particle in terms of material properties and geometry.

$$(P_p)_{\max} = 2\pi V_p^2 f \epsilon_o \kappa \frac{A_p}{(L_b - L_p)} \quad (9)$$

Now we make a gross approximation. We estimate that the thickness of the painted film is controlled to within a tolerance of 10%. This means that, if we try to make the film thickness equal to the particle thickness, to try to get maximum power transfer, then in the worst case,

$$L_b - L_p = \frac{1}{10} L_p \quad (10)$$

We can now combine equations (9) & (10) to get a simple estimate of the power transfer efficiency to a paint particle by capacitive means.

$$(P_p)_{\max} = (20\pi) \cdot (\epsilon_o \kappa) \cdot V_p^2 \cdot f \cdot L_p \quad (11)$$

Note the linear scaling of power with particle dimension. As the length scale gets smaller, the length to volume ratio goes up. This means that capacitive power transfer becomes capable of supplying a larger amount of power, relative to photovoltaic cells or internal energy storage, as the particle size considered becomes smaller.

For a 1 mm particle, a paint binder with a relative dielectric constant of 9.0, a resonant frequency of 1 MHz, and a semiconductor process with $V_{DD} = 3.3$ V, $P_{\max} = 100\mu\text{W}$. If we use a higher voltage process, such as HVCMOS, with $V_{DD} = 35\text{V}$, then $P_{\max} = 12$ mW. In this case, an on-particle switching converter [58] or linear regulator would be required to power digital circuitry at a lower voltage.

For our 110 μm paintable display particles, which consume $50\mu\text{W}$, we would need $f = 8$ MHz to transfer the required power with $V_{DD} = 3.3$ V. Alternately, we could make $V_{DD} = 10\text{V}$ (potentially by stacking some LED's in series) and use $f = 1.0$ MHz.

The efficiency of a power transfer system is defined to be

$$\eta = \frac{E_{\text{delivered}}}{E_{\text{delivered}} + E_{\text{lost}}} \quad (12)$$

where η is the efficiency, $E_{\text{delivered}}$ is the energy delivered to the load, and E_{lost} is the energy lost to the surroundings.

The energy lost per cycle in an RLC network is given by the Q of the network,

$$Q = \frac{E_{\text{stored}}}{E_{\text{lost}}} \quad (13)$$

where E_{stored} is the energy stored per cycle.

Since the magnitude of the impedance of the series capacitance and the particle are equal, and since they are in series with each other, we have,

$$(E_{\text{stored}})_{\text{series}} = E_{\text{delivered}} \quad (14)$$

The energy stored in the parallel capacitance, covering the parts of the paintable system where there are no particles, is given by

$$(E_{stored})_{parallel} = \frac{1}{2} C_{parallel} V^2 \quad (15)$$

where $C_{parallel}$ is the energy stored in the parallel capacitance per cycle, and V is the voltage across the capacitor plates. Since the impedance of the series capacitance and particle are matched,

$$V = 2V_p \quad (16)$$

We can write an expression for the parallel capacitance in terms of the paint area without a particle, A_b ,

$$C_{parallel} = \epsilon\kappa \frac{A_b}{L_b} \quad (17)$$

Assuming cubic particles with length L_p , that $L_b \approx L_p$, and given an area fill-factor of F for particles on the surface, we can write,

$$C_{parallel} = \epsilon\kappa \frac{L_p^2 \left(\frac{1}{F} - 1 \right)}{L_b} \approx \epsilon\kappa L_p \left(\frac{1}{F} - 1 \right) \quad (18)$$

We can find the energy stored in this capacitance by combining equations (15) and (18)

$$(E_{stored})_{parallel} = 2\epsilon\kappa L_p V_p^2 \left(\frac{1}{F} - 1 \right) \quad (19)$$

Working from equation (11), the energy delivered in a single cycle is

$$E_{delivered} = 20\pi\epsilon_o\kappa L_p V_p^2 \quad (20)$$

Combining equations (14), (19), and (20) we can write the total energy stored per cycle.

$$E_{stored} = \epsilon_o\kappa L_p V_p^2 \left(20\pi + \frac{2}{F} - 2 \right) \quad (21)$$

Combining equations (20) and (21), we get the ratio of energy stored to energy delivered.

$$\frac{E_{stored}}{E_{delivered}} = \frac{20\pi + \frac{2}{F} - 2}{20\pi} = 1 + \frac{\left(\frac{1}{F} - 1 \right)}{10\pi} \quad (22)$$

Combining equations (13) and (22), we get the ratio of energy lost to energy delivered.

$$\frac{E_{lost}}{E_{delivered}} = \frac{E_{stored}}{E_{delivered}} \cdot \frac{E_{lost}}{E_{stored}} = \left[1 + \frac{\left(\frac{1}{F} - 1 \right)}{10\pi} \right] \frac{1}{Q} \quad (23)$$

Combining equations (12) and (23), at long last, we get an expression for the efficiency.

$$\eta = \frac{10\pi Q}{10\pi(1+Q) + \frac{1}{F} - 1} \quad (24)$$

This is the expression for the efficiency of capacitive power transfer to paint particles, assuming a large, flat paint layer with a thickness tolerance of 10%, sandwiched inside of a large parallel-plate capacitor with a transparent front electrode. F is the area fill factor of paint particles, and Q is the quality factor of the LC network, before the addition of particles.

We can estimate Q at 40; this is a reasonable guess from electronics bench practice. The paintable display application detailed at the end of this report has an area fill factor of about 5%. From Equation 23, assuming perfect impedance matching and operation at resonance, capacitive power transfer has an efficiency of **96%**, similar to a good switching power supply. To get this kind of efficiency in practice, a feedback control system would probably be needed to keep the system precisely at resonance over variations in load.

2.3 Communications

2.3.1 Introduction

In this section, I compare physical mechanisms for implementing communication within a neighborhood of particles, to implement the model of inter-particle communication described by Butera in “Programming a Paintable Computer.” [13]

The paintable computing architecture assumes local communication; each particle can communicate with about 10 other particles in its local neighborhood; typically over a distance of less than 1 cm. Messages intended for particles further away need to be forwarded by the network.

This architecture is scalable to networks with millions of nodes; an architecture in which any particle could communicate with any other via a high-power transmitter would not be scalable, and would have higher total power consumption. There are three main reasons why local communication is scalable.

First, it reduces operating system overhead. If a particle had thousands of nodes inside its communications radius, it would be constantly processing and discarding messages not intended for it, and this would use processing resources and power.

Second, it reduces clutter on the communications channel, increasing the total available bandwidth. With large communications radiuses, the channel would always be jammed up with communication between far-off particles, so the data rate between any two particles would slow to a crawl.

Third, because of the inverse-square attenuation of electromagnetic waves, the required total transmit power goes as the square of the distance between nodes for direct communication,

but linearly with that distance for multihop communications. In the large distance limit, it takes lower total power to forward a message through a multihop network than to send it directly. [59]

In the following three sections, we discuss free-space optical, propagating-wave RF communications, and near-field RF communications.

Near-field RF communication has a very precise radius, and is not subject to multipath or fading, so it can be used for accurate particle localization. Optical communication has the lowest minimum die area requirement, and low peak power consumption. Optical communication is strictly line of sight. Both optical and near-field RF communications allow power efficiency as high as 30 pJ/bit.

Propagating-wave RF communication may be possible from volumes as small as 1 mm^3 , due to recent advances in antenna design and circuit techniques at 60 GHz. Propagating wave communication would be most useful between nodes with a very low fill factor; for example, nodes separated from one another by several meters.

There has been a recent revolution in RF CMOS circuit techniques. Except for very high performance systems, III-V nitride semiconductors are no longer required for RF. [60, 61, 62]

2.3.2 Communications Transports

2.3.2.1 Optical

Optical emitters and detectors (IR) can be fabricated on silicon, and the peak receiver power level required is very low. Optical emitters are directional. Diffusing lenses or a translucent, milky paint binder would be needed to insure communication between neighbors.

The intensity of an optical signal drops off with the inverse square of distance. The intensity of the emitter can be controlled by current regulation, starting low and ramping up, to establish a properly-sized neighborhood.

The circuitry and detectors for optical communications systems can be very small, and typically do not require external passive components. Workers at U.C. Berkeley built a “smart pixel” integrated optical receiver inside a $150 \text{ }\mu\text{m}$ square, using a 350 nm process. The data rate was 2.5 Mb/sec, and the receiver used only $50 \text{ }\mu\text{W}$ to achieve a -51dbm sensitivity. [63]

To produce -51 dBm at 10 mm from a 1 mm source requires -31 dBm of optical power. Given the ~10% efficiency of LED's, this requires -21 dBm of electrical power, about $8 \mu\text{W}$.

Using the figures listed above, for 1 mm nodes and 10 mm neighborhoods, we can estimate the total power consumption for the communications system at about $60 \text{ }\mu\text{W}$. This corresponds to an efficiency of **24 pJ/bit**.

2.3.2.2 Near-Field RF Communication

When particles are spaced by less than a few wavelengths, near-field communication becomes possible. Near-field communication works by capacitive or inductive coupling, rather than by propagation of electromagnetic waves.

Coils or capacitor plates for near-field communication can be much smaller than antennas at a given frequency: this allows the use of lower-frequency bands and less exotic circuit techniques than with propagating-wave communication.

Power drops off with the sixth power of distance. This gives brick-wall neighborhoods, makes eavesdropping and jamming all but impossible, and results in very low interference levels. However, it also restricts this technique to systems where the neighborhood size is physically small.

Nearfield systems are preferable for localization based on received signal strength, because, unlike propagating-wave systems, they are not subject to multipath or fading. In a near-field system, the received signal strength is a function only of the transmit power, the separation distance, and the presence of any dielectric (in the case of electric-field systems) or magnetic (in the case of magnetic-field systems) materials inside the neighborhood. Because magnetic objects (e.g. steel bolts) are rarer in nature than dielectric objects (e.g. raindrops, insects, people), magnetic-field systems are preferable for precise localization.

2.3.2.2.1 Near-Field Inductive Communications System Design

In this section, we present a block diagram and performance calculations for an inductive communications system. This system is designed for 1 mm particles and a 10 mm neighborhood size.

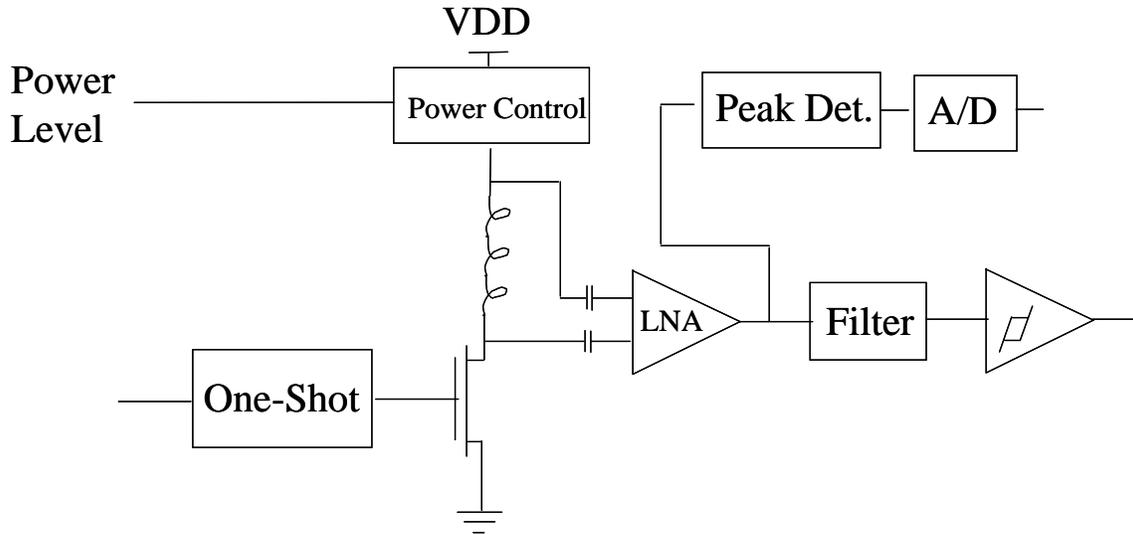


Figure 27: Block Diagram: Inductive Communications System: This system includes power control and RSSI, for precise control of neighborhood size, precise localization, and very high data rate.

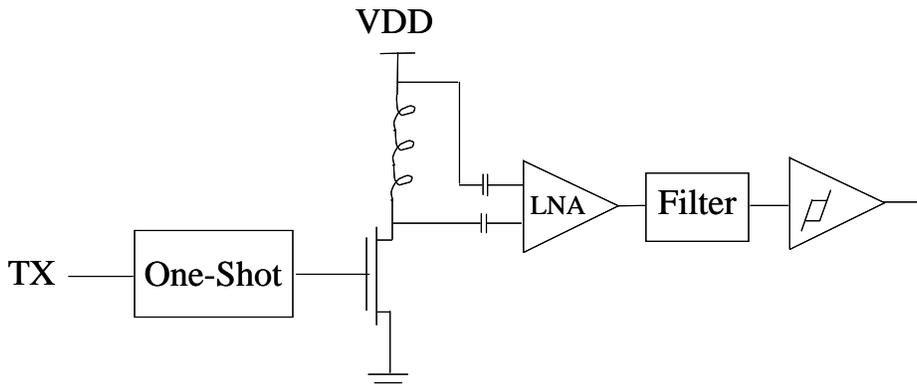


Figure 28: Block Diagram: Inductive Communications System: This system aims for minimum transistor count. All of the blocks shown are simple analog elements and require just a few transistors each; the total transistor count for this circuit is probably 20-30.

In this section, we will calculate the resistance, self-inductance, and mutual inductance of the inductors that couple data between particles. We will use these figures to calculate the path loss at 10 mm, signal to noise ratio, bandwidth, transmitter and receiver power consumption, and maximum data rate.

For this design example, we select an Intel 130 nm CMOS process, and a die size of 1mm². We place the communications inductor on metal layer 6 (see Figure 29 for design rules)

and select $d_{in} = 0.8$ mm and $d_{out} = 1$ mm. (See Figure 30) We select an analysis frequency of 1 GHz, corresponding to the transmission of 1 ns pulses.

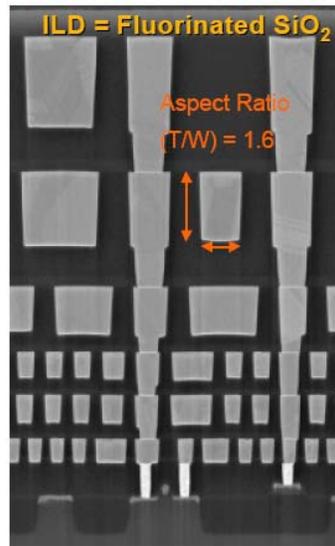


Figure 29: Scanning Electron Micrograph of a 130nm CMOS IC, Cross Section: The metallization pitch is 350 nm for the bottom metal layer and 1200 nm for the top metal layer, number 6. Upper layer traces are 600 nm wide by 960 nm high. From [71].

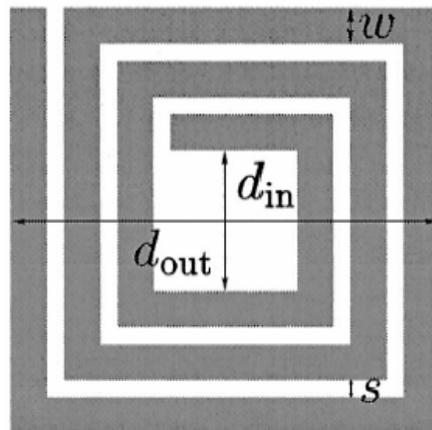


Figure 30: Planar Integrated Circuit Inductor from [64]

First, we need to compute the diameter and area fill factor of the inductor.

$$d_{avg} = \frac{d_{in} + d_{out}}{2} \quad (\text{from [64]}) \quad (25)$$

$$\rho_{fill} = \frac{d_{out} - d_{in}}{d_{out} + d_{in}} \quad (\text{from [64]}) \quad (26)$$

For this design, using the above expressions, $d_{avg} = 0.9$ mm and $\rho_{fill} = 0.1111$. Given small ρ_{fill} and thin traces, so that the skin effect is not an issue, we can write

$$R = \rho \frac{L}{A} \cong \rho \frac{4d_{avg}N}{W_{turn}H_{turn}} \quad (27)$$

In this equation, ρ is the resistivity of the inductor trace, L is the total trace length, A is the cross-sectional area of the trace, N is the number of turns, W_{turn} is the width of the trace, and H_{turn} is the height of the trace.

We select $W_{turn} = 8$ μm and set $H_{turn} = 960$ nm, corresponding to the value given in Figure 29. Given these trace dimensions, we can fit about 6 turns, so $N = 6$. The traces are made of copper, so we set $\rho = 1.6 \times 10^{-8}$ Ω/m .

From equation (27), we get $R = 49.2$ Ω . This is not an accident; we selected the trace width to make the resistive part of the impedance of the inductor as close as possible to the standard 50 Ω .

A standard expression for the inductance of a planar spiral is,

$$L = 2.34\mu_o \frac{N^2 d_{avg}}{1 + 2.75\rho_{fill}} \quad (\text{from [64]}) \quad (28)$$

where μ_o is the magnetic permeability of free space, equal to 1.26×10^{-6} H/m. From this expression, we see that the inductance of our inductor is 79 nH.

We can calculate the reactive part of the impedance of the inductor using

$$|X_L| = 2\pi fL \quad (29)$$

The reactive part of the impedance of the inductor at 1 GHz is 497 Ω . This is good news: it means the inductor looks 10 times more like an inductor than a resistor at 1 GHz.

The mutual inductance between two identical inductors oriented in the same direction is given by

$$L_M = \frac{\mu_o \pi N^2 (d_{avg}/2)^4}{2r^3} \quad \text{from [65]} \quad (30)$$

where L_M is the mutual inductance and r is the separation distance between the inductors. Taking $r = 10$ mm, equal to the neighborhood size, $L_M = 3.2$ pH.

We can compute the induced voltage by combining the expressions for self and mutual inductance. When the system transmits a pulse, V_{DD} is forced across the transmitting inductor.

This induces a current in the transmitting inductor according to

$$V_{DD} = L \frac{dI_T}{dt} \quad (31)$$

where I_T is the current in the transmit inductor and t is time. This time varying current produces a time-varying magnetic field, which induces a voltage in the receive inductor according to

$$V_r = L_M \frac{dI_T}{dt} \quad (32)$$

where V_r is the voltage induced in the receive inductor. Combining equations (31) and (32), we get an expression for the voltage of the received signal in terms of V_{DD} ,

$$V_r = \frac{L_M}{L} V_{DD} \quad (33)$$

For a design using the Intel 130 nm process, $V_{DD} = 1.2V$, so $V_r = 47 \mu V$. This voltage appears in series with the 50Ω resistive impedance of the inductor, so we can compute the received power using

$$P_{signal}(dbm) = 10 \log_{10} \left(\frac{V_r^2}{(50\Omega)(1mW)} \right) \quad (34)$$

where $P_{signal}(db)$ is the signal power at the receiver, in decibels relative to 1 mW. Using this expression, we calculate that the signal power at the receiver is -73 dBm.

We can calculate the noise power using

$$P_{noise}(dbm) = 10 \log_{10} \left(\frac{kTf}{Q(1mW)} \right) + NF(db) \quad (35)$$

where $P_{noise}(dbm)$ is the input-referred noise power at the receiver, Q is the quality factor of the receive amplifier, related to the half-power bandwidth, and $NF(db)$ is the noise figure of the receive amplifier in decibels.

Selecting a receive amplifier with $Q = 10$ (for a bandwidth of 100 MHz at $f = 1$ GHz), and a noise figure of 3 dB, both of which are readily achievable using an inductively-degenerated CMOS LNA [62], we can use equation (35) to calculate that the input-referred noise at the receiver is about -90 dbm. This gives a signal-to-noise ratio of 17 dB, about 50 expressed as a straight power ratio.

We can estimate the communications bandwidth for the channel using the Shannon capacity. At one time, this would have been a gross over-estimation, but modern coding techniques (e.g. Turbo codes) can come very close to saturating the Shannon capacity.

$$C = W \log_2(1 + SNR) \quad (36)$$

Substituting $W = 100$ MHz and $SNR = 50$, we get a channel capacity of 567 Mb/s.

We can get bitrates close to the channel capacity using 8-QAM modulation and turbo coding. Alternatively, we can use OOK modulation, as is shown in the block diagrams, and accept lower data rate, probably around 50 Mb/s, in exchange for reduced area.

Based on designs given in [62], we can guess that a 1 GHz, $Q = 10$ inductively degenerated LNA with 3 dB NF will require about 10 mW. This is likely to be the dominant source of receiver power; we estimate the receiver power at 1.5 times this figure, which is 15 mW.

We can estimate the transmitter power using

$$P_T = 2 \frac{V^2}{|R + jX_L|} \quad (37)$$

which assumes that the power stored in the inductor is completely dissipated on every cycle, and Class-A (worst case) amplification. It may be possible to do much better than this with Class-D amplification and power recovery from the inductor. From Equation (37), we can estimate the transmitter power at 6 mW.

$$\left(\frac{E}{bit} \right) = \frac{P_D}{C} \quad (38)$$

Combining all of our assumptions and approximations, some liberal and some conservative, we can estimate power efficiency of this communications system, to a node at the edge of the neighborhood, at **37 pJ/bit**.

2.3.2.3 Propagating-Wave RF Communication

The most obvious challenge to building a propagating-wave RF transceiver inside of 1 mm^3 is the size of the antenna.

Commercial RF systems use external antennas, so the transmitter output and receiver input are almost always matched to the industry-standard 50Ω impedance. The radiation resistance of an antenna is related to its length; for example a 75Ω antenna must have an electrical length of about $1/2$ wavelength. [66] A 1-mm long, $1/2$ wavelength antenna would have a center frequency of 150 GHz. The design of electrically shorter antennas and the corresponding low impedance RF circuits to drive them is possible in principle. Of course, the design of higher frequency RF circuits, still operating at the same comfortable impedance, is also possible in principle.

There is a recent body of work concerning techniques for fabricating integrated antennas. In one result, workers fabricated a 2-mm long zigzag dipole antenna, and observed proper inverse-square dependence over 4-5 meters, at a frequency of 24 GHz. [67]

A great deal of work is ongoing to develop circuit techniques for low-power RF communications. For example, workers fabricated a complete 433 MHz UHF radio transceiver, with 24 kb/s data rate and 1 mW power consumption. [68]

The FCC recently created a new ISM band at 60 GHz, with 7 GHz bandwidth. As a result, the literature is filled with successively lower power and lower area radios, some fabricated using CMOS technology, for high-rate communication at 60 GHz over distances of about 10 m. [69]

Given the zig-zag dipole result mentioned earlier, it seems clear that a 60 GHz antenna could fit within 1 mm^3 . Given the rate at which 60 GHz transceivers are shrinking, due to the intense interest in 7 GHz of unlicensed spectrum, it seems likely that in the coming years we will see a complete 60 GHz radio, including antenna, that can fit inside 1 mm^3 . This kind of radio would be useful for systems with very sparse nodes, separated by many meters.

2.4 CPU and Memory

The technical motivation for building millimeter-sized computing systems comes from the very small amount of silicon area required by the CPU and memory. Still, there are limits on the size of CPU and especially the size of a memory array that can fit within a given area given state-of-the-art lithography, so in this section, we state these limits for reference.

2.4.1 CPU

The centimeter-scale paintable display described in the application examples uses the Atmel ATFR40162 processor, which has an ARM7TDMI core. This is a 32-bit RISC integer machine. On a 130 nm process, the ARM7TDMI fits inside 0.26 mm^2 of silicon area, and consumes **60 $\mu\text{W}/\text{MHz}$** , with an instruction rate of 0.9 MIPS/MHz, and runs at up to 133 MHz. [70] This core is fully static, so it can be run at lower speeds with a proportional decrease in power.

2.4.2 RAM

A static RAM cell has six transistors, consumes **0.57 μm^2** of silicon area on a 65 nm process [11], and $2.0 \mu\text{m}^2$ of silicon area on a 130 nm process. [71] The 256K x 8 SRAM array in the V1.0 paint particle requires about 2×10^6 cells, for a total area of 1.1 mm^2 on a 65 nm process and 4 mm^2 on a 130 nm process.

If a larger amount of RAM is needed DRAM could be used, but at the cost of static power consumption. A one-bit DRAM cell requires **0.11 μm^2** of silicon area on a 65 nm process. [72]

For the construction of machine with a reasonable amount of RAM, RAM area is much more important than processor area. For a 1 mm^2 computer fabricated on a 65 nm process, a 32-bit processor will easily fit, but only 64K x 8 of SRAM or 320K x 8 of DRAM will fit.

2.4.3 FLASH

FLASH memory (or mask ROM) will be required for operating system storage by paint particles, and may also be useful for long-term zero-power storage of information for applications.

FLASH memory requires much less area per bit than SRAM; **0.16 μm^2** per cell on a 130 nm process. [73] A 256K x 8 FLASH memory array requires 0.32 mm^2 . While the industry is confident that FLASH memory will eventually scale down to 45 nm processes and beyond, [74] at the time of writing, this appears to be the smallest published FLASH cell.

2.5 Light Emitting Diodes

2.5.1 Power Requirements

In this section, we determine how much electrical power is needed, per particle, for a color display with a particular brightness in Nits. Typical display brightness values are 200 Nits for indoor-readable LCD monitors and 1500 Nits for a sunlight-readable displays.

To get information on the efficiency of light emitting diodes, we obtained datasheets from Lumileds Corp. on high-efficiency red, green, and blue LED's available in bare die form. This data is reprinted below.

Color	Process	Sample Device	Data
Red 626 nm	TSAlInGaP	HWFR-B517 (Lumileds)	0.42 lm @ 40mA, 2.3 V
Green 520 nm	InGaN	HWFR-P5G2 (Lumileds)	4.4 lm @ 50 mA, 4.0 V 2.2 lm @ 20 mA, 3.2 V
Blue 475 nm	InGaN	HWFR-P5B2 (Lumileds)	1.6 lm @ 50 mA, 4.0 V 0.8 lm @ 20 mA, 3.4 V

Table 4: *Electro-Optical Efficiency Information from LED Data Sheets*

According to the Lumileds InGaN data sheet, bare die LED's emit light over 2.6 steradians of solid angle. From this information, we calculated the luminous intensity of each LED in candela. Also, by multiplying the stated voltage times the stated current, we calculated the electrical power consumed at the given conditions.

Color	Luminous Flux at Test Condition	Luminous Intensity at Test Condition	Electrical Power at Test Condition	Efficiency at Test Condition
Red 150 lm / W	0.42 lm	160 mcd	92 mW	3.0%
Green 500 lm/W	2.2 lm	846 mcd	64 mW	6.9%
Blue 100 lm/W	0.8 lm	308 mcd	68 mW	11.8%

Table 5: Calculated Electro-Optical Efficiency Information

Color	Electrical Power per Luminous Intensity
Red	588 mW / candela
Green	76 mW / candela
Blue	222 mW / candela
Total (RGB White)	886 mW / candela

Table 6: Electro-Optical Efficiency of RGB LED's

Based on this information, we developed the following formula to give the power requirements for an LED display made with red, green, and blue LED's:

$$P = \left(0.886 \frac{W}{candela} \right) \cdot (A \ m^2) \cdot \left(B \frac{candela}{m^2} \right) \quad (39)$$

where B is the brightness of the display in Nits (candela/m²), and A is the area of the display in m².

From Table 5, we can see that the efficiency of LED's is between 3% and 12%, so most of the power predicted by this formula is dissipated as heat, and only a small fraction leaves the particle as photon flux. This means that heat dissipation can place a limit on the minimum particle size for a display of a given brightness and particle spacing.

Display	Power Required
LED monitor 30 cm square box 200 nits brightness	17 W
Sunlight Readable Full-Color Paintable Display (per node) 0.25 mm ² display area per particle 1500 nits brightness	325 μW
Indoor Readable Full-Color Paintable Display (per node) 0.25 mm ² display area per particle 200 nits brightness	43 μW

Table 7: LED Power Requirements

2.6 Actuation

2.6.1 Why Distributed Action?

In our design study of millimeter-scale programmable matter, we chose to examine actuation mechanisms without moving parts in the particles. The particles we describe and analyze contain electrodes, wires, or permanent magnets which create electric or magnetic fields around the particle. Particles cooperate with one another to apply forces and cause relative motion. We call this approach “distributed actuation.”

It might be possible to actuate programmable matter using clockwork and traditional machine elements (gears, levers, pulleys, belts, wheels) but there are several reasons to consider distributed actuation instead.

First, fabrication of sub-millimeter mechanical elements is non-trivial, requiring a large number of masks and time-consuming, expensive processes. In contrast, the electrodes for the distributed electrostatic actuators we propose can be built on the metal layers of standard CMOS.

Second, friction and stiction are serious issues in the design of a system with sub-millimeter moving parts, because surface effects become relatively stronger with decreasing length scale. A geared actuation system would require mechanical parts much smaller than the particles. In contrast, using distributed actuation, the smallest moving parts are the particles themselves. It is straightforward to demonstrate that stiction is not stronger than gravity for dry millimeter-sized particles of silicon --- dry sand doesn’t stick to the ceiling.

Finally, sub-millimeter protrusions from particles (e.g. wheels) could be prone to breakage. MEMS devices are typically packaged, so none of the internal parts are exposed. Since programmable matter particles are intended to operate in a natural environment, it is

important to make the nodes mechanically robust. Not having external protrusions means that the nodes can be smooth hard-units, impervious to mechanical damage except by crushing or extreme wear.

2.6.2 Electrostatic vs. Magnetic

Distributed magnetic actuation is possible at the centimeter scale, but is inefficient due to the high resistive losses inherent in magnetic actuators designed to produce large forces at low speed. Magnetic actuation becomes even more inefficient as the size scale is reduced, unless permanent-magnet actuator designs are used. However, incidental, undesired forces between permanent magnets are a serious problem at centimeter scale, and become even more serious as size scale is reduced. Because of this, we are not optimistic about the prospects for a feasible millimeter-scale distributed magnetic actuation system.

It is important to note the above conclusion is contingent upon the conductivity of the available materials. If a material with conductivity 10-100 times higher than copper were to become available (e.g. a room-temperature superconductor) magnetic actuation for millimeter-scale programmable matter would be more practical.

Distributed electrostatic actuation is awkward at the centimeter scale, due to the requirement for high voltage. The large electronic components required to switch the required voltage and the need for a liquid dielectric to prevent air breakdown limit the system to operation in an enclosed, specially prepared environment.

At millimeter scales, these concerns are eliminated, and distributed electrostatic actuation becomes extremely attractive.

2.6.3 Magnetic Field Systems

2.6.3.1 The case against magnetically actuated “Utility Fog”

If it were feasible, a particularly attractive system design would be the “Utility Fog” concept from science fiction where particles occupy random positions in space, not necessarily in contact, suspended against gravity by inter-particle magnetic fields. In this section, we show that millimeter-scale, magnetically-actuated utility fog is not feasible due to its high power dissipation. This analysis is included because it motivates the study of the power-efficient but more complicated motion systems in subsequent sections.

For generic motion, particles cannot contain soft magnetic material, because soft magnetic material is attracted to magnetic fields. Soft magnetic material is heavily used in motors due to its ability to guide flux and decrease the effective magnetic path length, but in a motor, the attraction between the rotor and stator is balanced by restoring forces from the bearing. In a

utility fog system there is no contact, in general, and no bearings or bearing surfaces. Thus, for generic, six-axis actuation without bearings, we need to restrict our analysis to nodes containing current-carrying wires only.

Without loss of generality, let's us consider nodes containing wire coils, arranged coaxially. This will allow us to compute the power required for (unstable) levitation of a node at a given height. Since coaxial levitation is a subset of the required set of capabilities, and coaxial coils are an ideal way to accomplish this particular subset, this will place a lower bound on the power required for generic actuation in a utility fog system.

The mutual inductance of two coaxial wire loops of radius r and axial separation distance x , from the Biot-Savart law, is

$$L(x) = \frac{\pi\mu_0}{2} \frac{r^4}{(r^2 + x^2)^{3/2}} \quad (40)$$

The mechanical force F along the direction of x between two inductors carrying current i with mutual inductance $L(x)$ is

$$F(x) = \frac{dW}{dx} = \frac{1}{2} \frac{dL(x)}{dx} i^2 \quad (41)$$

Combining Equations (40) and (41), the force between two wire loops as a function of their separation distance x is

$$F_{loop}(x) = \frac{3\pi\mu_0}{4} \frac{r^4 x}{(r^2 + x^2)^{5/2}} i^2 \quad (42)$$

Now, considering two coaxial solenoids of length L , spaced an on-axis distance s , each with a large number of turns N , the force between them is

$$F(x) = \frac{N}{L} \int_0^L \frac{N}{L} \int_0^L F_{loop}(s + x_1 + x_2) dx_2 dx_1 \quad (43)$$

Performing the elementary integration in equation 4, we arrive at Equation 5, the expression for the force between coaxial solenoids.

$$F(x) = \frac{\pi}{4} \mu_0 N^2 i^2 \left(\frac{r}{L} \right)^2 \left[\frac{2(L+s)}{\sqrt{r^2 + (L+s)^2}} - \frac{2L+s}{\sqrt{r^2 + (2L+s)^2}} - \frac{s}{\sqrt{r^2 + s^2}} \right] \quad (44)$$

In equation 5, the terms in brackets and parentheses are scale-invariant --- the force depends only on the current and number of turns.

We will now compute the power dissipation in such a coil at stall, due to resistive losses in the conductor. For a winding thickness t , where $t \ll r$, and a conductor of resistivity ρ , the resistance is

$$R = \frac{(Length)}{\sigma(Area)} = \frac{2\pi r N^2}{\sigma L} \quad (45)$$

So the power dissipation in the two coils, each with resistance R, due to the current i, is

$$P = 2i^2 R = \frac{4\pi r N^2 i^2}{\sigma L} \quad (46)$$

Combining Equations (44) and (46), we obtain the expression for the F(x) as a function of resistive loss P.

$$F(x) = \frac{\pi}{4} \mu_0 \sigma P \frac{tr}{L} \left[\frac{2(L+s)}{\sqrt{r^2 + (L+s)^2}} - \frac{2L+s}{\sqrt{r^2 + (2L+s)^2}} - \frac{s}{\sqrt{r^2 + s^2}} \right] \quad (47)$$

Neglecting obvious issues of controllability for the moment, the force required to levitate an object against gravity is mg. Assuming cubic nodes of side length D, with mass density ρ

$$F(x) = \rho D^3 g \quad (48)$$

So we can approximate the power per node required for actuation in a magnetic levitation-based system by combining equations (46) and (47).

$$P = \frac{4}{\pi} \left(\frac{\rho g}{\mu_0 \sigma} \right) \frac{DL^3}{tr} \left[\frac{2(L+s)}{\sqrt{r^2 + (L+s)^2}} - \frac{2L+s}{\sqrt{r^2 + (2L+s)^2}} - \frac{s}{\sqrt{r^2 + s^2}} \right]^{-1} \quad (49)$$

The power available for actuation is limited in three ways, which we will now introduce as constraints. First, it is limited by the volumetric power density of on-board power supply, which we will call χ , with units of W/m³.

$$P < \chi D^3 \quad (50)$$

Second, it is limited by the allowable heat dissipation from each node. The heat transfer \dot{Q} from a node levitating in free space is the sum of the heat transfer due to conduction, \dot{Q}_c and the heat transfer due to radiation \dot{Q}_r .

$$P < \dot{Q}_c + \dot{Q}_r \quad (51)$$

The heat transfer due to radiation is given by the Stefan-Boltzmann law, where T is the operating temperature of the node, σ is the Stefan-Boltzmann constant, $5.670 \times 10^{-8} \text{ JK}^{-4} \text{ m}^{-2} \text{ s}^{-1}$, and ϵ is the emissivity of the node, of which about 0.8 is typical for non-reflective engineering materials.

$$\dot{Q}_r = \epsilon A \sigma (T^4 - T_{amb}^4) \quad (52)$$

The heat transfer due to convection can be approximated using the Nusselt number correlation for natural convection on a sphere. [3]

$$Nu_D = 2 + \frac{0.589 \cdot Ra_D^{1/4}}{\left[1 + (0.469 / Pr)^{9/16} \right]^{4/9}}; Ra \leq 10^{11}; Pr > 0.5 \quad (53)$$

The Prandtl number Pr for air is 0.71, so that constraint is satisfied. The Raleigh Number Ra is given in terms of the Grashof number Gr as

$$Ra_D = \frac{(bouyancy)}{(viscosity)} = Gr_D \cdot Pr \quad (54)$$

The Grashof number is given by

$$Gr_D = \frac{(bouyancy)(inertia)}{(viscosity)^2} = \frac{\beta(T - T_{amb})gD^3}{\nu^2} \quad (55)$$

Where g, D, T, and T_{amb} are as previously defined, ν is the kinematic viscosity of air, 2×10^{-5} m²/s. β , the coefficient of thermal expansion of air, can be written in terms of the film temperature in the natural convection region as

$$\beta = \frac{1}{T_{film}} \quad (56)$$

We can approximate the film temperature as the average of the ambient temperature and the node temperature

$$T_{film} = \frac{T + T_{amb}}{2} \quad (57)$$

Finally, so long as equation (53) is valid, we can write the heat transfer due to natural convection, here assuming a spherical node for simplicity, as

$$\dot{Q}_c = Nu_D k \pi D (T - T_{amb}) \quad (58)$$

Where k is the thermal conductivity of air, about 0.26 W/m-K.

Finally, the steady-state power for actuation is limited by the volumetric energy density ξ (J/m³) of the supply, and by the required run-time τ .

$$P < \frac{\xi D^3}{\tau} \quad (59)$$

The above analysis is summarized in Figure 31, for some aggressive numerical values for material properties and operating conditions: T = 125°C, the maximum junction temperature of semiconductors, and T_{amb} = 25°C. The conductivity of copper is used for σ , 5.96×10^7 S/m. The power source is a lithium-polymer battery, with $\xi = 1.1 \times 10^9$ J/m³ and $\chi = 2.8 \times 10^6$ W/m³. We take L=D, and r=D/2, assuming that the coil completely fills the node, and t = D/10, since the analysis is only valid for small t, and any real system would need at least three orthogonal coils and room for the battery.

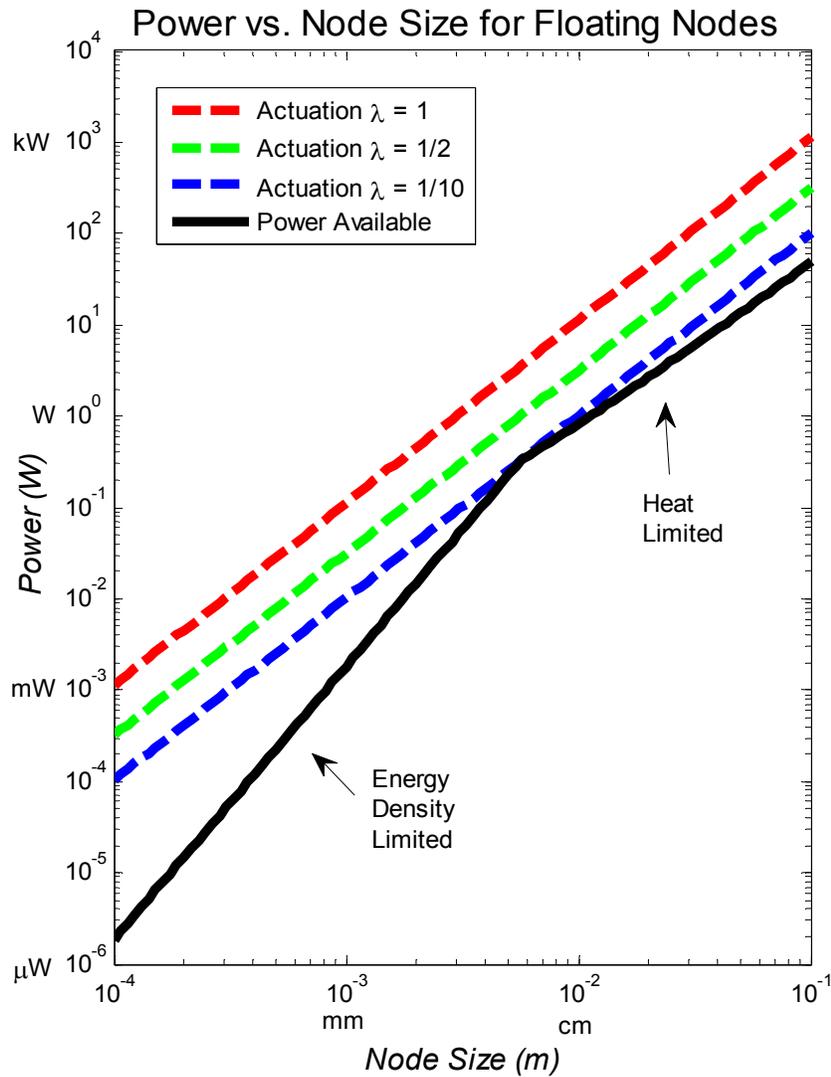


Figure 31: Best-case values of material properties and operating conditions still show infeasibility for millimeter-scale magnetic utility fog. λ is the ratio of levitation distance to node size. Nodes dissipate too much power at every size scale. The source of the power dissipation is resistive losses in the field-producing coils. Above 4 mm power dissipation is limited by heat dissipation from natural convection, and below 4 mm, power dissipation is limited by the energy density of lithium polymer batteries inside the nodes.

2.6.3.2 Permanent-Magnet Linear Motors

A current-carrying wire in the magnetic field of a permanent magnet experiences a force. This is the basic principle of operation of small electric motors. The spirit of this analysis is generality, so we encapsulate out almost everything specific to a particular geometry into a dimensionless parameter, so we can instead concentrate on asymptotic performance and scaling.

Direct-drive distributed actuation of programmable matter using permanent magnet actuators is more feasible than the approach described in the previous section, but still has low efficiency compared to traditional electric machines. This is because the efficiency of magnetic motors is a strong function of the relative speed between the two halves of the motor. The efficiency of a permanent-magnet motor has a peak value at a particular operating speed, and goes to zero at either very low or very high speeds. If efficient electrical-to-mechanical energy conversion is required at some speed other than this speed, power transmission utilizing mechanical advantage (e.g. gears, levers, belts, screws) is required. To design programmable matter particles without internal moving parts, we cannot employ these mechanisms.¹

The efficiency of an electric motor is the ratio of output power to input power. The output power is the mechanical power at the shaft, force times velocity. The input power is the electrical power at the electrical terminals, voltage times current.

$$\eta = \frac{P_{out}}{P_{in}} = \frac{Fv}{VI} \quad (60)$$

Invoking the first law of thermodynamics with a control volume around the entire motor, we see that the electrical power into the motor equals the mechanical power output plus the heat dissipation Q .

$$VI = Fv + Q \quad (61)$$

The loss mechanism most important for low speed operation is resistive loss in the windings. At high speeds, eddy currents induced into the magnetic materials, kinetic friction in the bearings, and aerodynamic drag are the primary loss mechanisms. For programmable matter, we are interested in operation at the low-speed limit, so in this analysis we will consider only the winding resistance R .

$$Q = I^2 R \quad (62)$$

¹ A large-scale programmable matter system without internal moving parts might be made to emulate machine elements to do large scale work, given an effective latching technology, but at the single-node scale, the actuator and load are unified and directly coupled.

The force produced by an electromagnetic system is given by the Lorentz force law, where \vec{f} is the force density, q is the charge density, \vec{E} is the electric field, \vec{J} is the current density, and \vec{B} is the magnetic flux density.

$$\vec{f} = q\vec{E} + \vec{J} \times \vec{B} \quad (63)$$

In a magnetic motor, the $q\vec{E}$ (electrostatic) term is negligibly small, so we restrict our attention to the $\vec{J} \times \vec{B}$ (magnetostatic) term.

$$\vec{f} = \vec{J} \times \vec{B} \quad (64)$$

Electric motor and actuator design is the art of designing the shape of the current density and magnetic field, the right hand terms of equation (64), to produce the desired force density (e.g. force or torque in the desired direction) on the left hand side. The performance of a particular geometry is often summarized via the force constant k_m . k_m encapsulates the geometry of the windings, strength of the permanent magnets, and flux-focusing action of any magnetic materials into one number.

$$F = k_m I \quad (65)$$

Substituting equations (60), (61), (62), and (65), we get an expression for the efficiency of a linear motor at low speed:

$$\eta = \frac{F \mathcal{G}}{F \mathcal{G} + F^2 \left(\frac{R}{k_m^2} \right)} \quad (66)$$

From this expression, one can see that the efficiency of a motor is zero at zero speed, and increases sub-linearly as the speed increases. The stall power is the electrical power input to the motor at zero speed. From Equation (66), we can see that:

$$P_{stall} = F^2 \left(\frac{R}{k_m^2} \right) \quad (67)$$

The ratio of the stall force output of a motor to the square root of the stall power input has a special name in the electric machine industry: it is called the “fundamental motor constant.”

$$\left(\frac{F}{\sqrt{P_{stall}}} \right) = F^2 \left(\frac{R}{k_m^2} \right) \quad (68)$$

To evaluate the feasibility of permanent-magnet programmable matter, we need to compare the power required for reconfiguration with the power available.

To do this, we will consider a generic, best-case permanent magnet DC motor, without specifying the geometry to achieve this best case, which will set an absolute lower limit on the fundamental motor constant.

Consider the armature winding of a motor, of length l , with permanent magnets with a remnant magnetization B_0 . If the motor is perfectly designed, so that the full remnant magnetization cuts the windings exactly perpendicularly, then, from the Lorentz force law, we can write

$$F = IlB_0 \quad (69)$$

Such a motor would have an armature resistance R of

$$R = \frac{l}{\sigma A} \quad (70)$$

Where the stall power is

$$P_{stall} = I^2 R \quad (71)$$

Combining equations (69), (70), and (71), we see that

$$\left(\frac{F}{\sqrt{P_{stall}}} \right) = B_0 \sqrt{Al\sigma} \quad (72)$$

Recognizing Al as the volume of the copper windings, and thus a minimum volume for the motor v , we get

$$\left(\frac{F}{\sqrt{P_{stall}}} \right) < B_0 \sqrt{v\sigma} \quad (73)$$

We now define a dimensionless version of the fundamental motor constant, α , to scale (73), a lower bound, to (74), which is an exact equality.

$$\left(\frac{F}{\sqrt{P_{stall}}} \right) = \alpha B_0 \sqrt{v\sigma} \quad (74)$$

$$0 \leq \alpha \leq 1$$

To summarize, F is the force output, P_{stall} is the electrical power input at stall, α is a constant from 0 to 1, B_0 is the remnant magnetization of the permanent magnet material, v is the volume of the motor, and σ is the electrical conductivity of the windings.

In reality, the flux and current are not perpendicular everywhere, and significant volume inside the motor is needed for the permanent magnets, iron components, bearings, and housing of the motor, so α is much less than 1.

Re-arranging Equation (75), substituting ρvg for F , we arrive at the power required to actuate a programmable matter node with particular dimensions and a particular dimensionless overall fundamental motor constant.

$$P = \frac{\rho^2 g^2}{\alpha B_0 \sigma} D^3 \quad (75)$$

Motor	Type	Dimensions	Fundamental Motor Constant	Alpha
Servotube 3804 (Copley Controls, Canton, MA)	Radial-Flux Tubular Linear Motor	70 x 84 x 218 mm	$14.54 \frac{N}{\sqrt{W}}$	0.04
HyCore-10 LMHS10A-3COA (Baldor Motion Products)	Toothed-Track Linear Motor	50 x 64 x 118 mm	$9.1 \frac{N}{\sqrt{W}}$	0.05

Table 8: Dimensionless Fundamental Motor Constant (Alpha) for some commercial linear motors.

Figure 32 shows the stall power (from Equation (75)) vs. scale, for several values of alpha, plotted on the same axes as the node power constraint worked out in §2.6.3.1

Because the stall power is proportional to the volume of the nodes, as is the power available from an internal battery, permanent magnet linear motors are scale invariant below 4 mm, where the power constraint switches from heat-limited to energy-density-limited.

From Figure 32, we see that permanent magnet actuation for programmable matter particles may be possible, but requires just slightly less than the power available, due to the need for six independent actuators, and the generally low volumetric efficiency of state-of-the-art permanent-magnet linear actuators versus the limit described above.

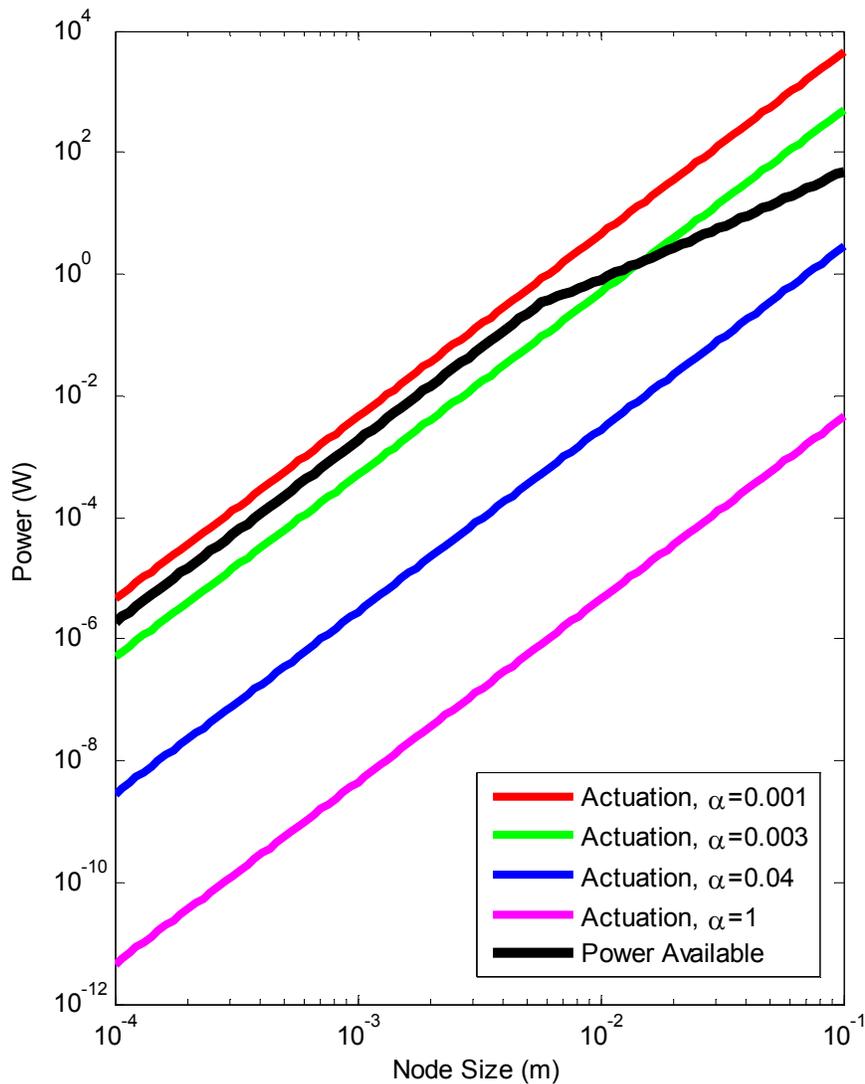


Figure 32: Power vs. Node Size for Permanent Magnet Linear Actuation. Magnetic actuation for programmable matter may be possible, but the appropriately scaled power consumption of best-practice commercial linear motors is very close to the maximum power available. A theoretical bound for linear motors made with copper and rare earth magnets (derived above) is at $\alpha = 1$, and commercial linear motors have $\alpha = 0.04$. Best-case programmable matter design is represented by $\alpha = 0.003$, for nodes with half their volume devoted to motors, and six independent axes of motion each requiring separate motors. Worst-case programmable matter design is represented by $\alpha = 0.001$, assuming motors three times worse than commercial.

2.6.3.3 Scaling Analysis of Permanent-Magnet Interactions

The force of attraction or repulsion between two permanent magnets becomes relatively stronger than weight or permanent-magnet linear motor action with decreasing length scale. This places a lower limit on the feasible length scale for a magnetic distributed actuation system based on nodes with unsecured, unshielded permanent magnets; as length scale decreases, eventually parasitic attraction between the permanent magnets on different nodes overwhelm any other behaviors we may try to engineer into the system.

In the remainder of this section we will work out the scaling law for the force between permanent magnets, to justify the statements in the previous paragraph.

High performance rare-earth magnets have the straight-line demagnetization relationship given below, where B is the magnetic flux density, H is the magnetic field strength, and B_0 is the remnant magnetization.

$$B = \mu_0 H + B_0 \quad (76)$$

Invoking the definition of magnetization,

$$\vec{B} = \mu_0 (\vec{H} + \vec{M}) \quad [75, \quad \text{p.349}] \quad (77)$$

and combining Equations (76) and (77), we see that the magnetization vector for a rare-earth magnet, magnetized in the k direction, is

$$\vec{M} = \hat{k} \frac{B_0}{\mu_0} \quad (78)$$

Working directly from Ampere's law, a magnetic material with magnetization M can be modeled as an equivalent current density J ,

$$\vec{J}_{eq} = \nabla \times \vec{M} \quad [75, \quad \text{p. 348}] \quad (79)$$

For a solid magnet with a uniform magnetization, the curl of the magnetic field is zero everywhere except at the surface, so the effect of the magnetization can be replaced by an equivalent current on the surface of the magnet.

For the case of an axis-aligned rectangular block magnet or disc magnet, the equivalent current flows in a flat loop around the x and y plane sides, with surface current density given by Equation (80).

$$J_{s,eq} = \frac{B_0}{\mu_0} \quad (80)$$

Without loss of generality, consider a thin disc-shaped permanent magnet, with radius R and thickness t . It appears as a current loop, with current

$$I_{eq} = \frac{B_0}{\mu_0} t \quad (81)$$

From Equation (79), recall that the force between two coaxial wire loops is

$$F_{loop}(x) = \frac{3\pi\mu_0}{4} \frac{r^4 x}{(r^2 + x^2)^{5/2}} i^2 \quad (82)$$

Combining Equations (81) and (82) gives the force between two coaxial disc-shaped thin permanent magnets.

$$F_{magnets}(x) = \left(\frac{3\pi B_0^2}{4\mu_0} \right) \frac{r^4 t^2 x}{(r^2 + x^2)^{5/2}} \quad (83)$$

Rewriting Equation (83) in terms of a non-dimensional function f and particle size D , we can better see the scaling behavior.

$$F_{magnets}(x) = \left(\frac{B_0^2}{\mu_0} \right) f\left(\frac{r}{D}, \frac{t}{D}, \frac{x}{D} \right) D^2 \quad (84)$$

The left-hand term in Equation 9 is a set of material properties and physical constants, so it is scale-invariant. The right-hand multiplicative term has the units of length squared, so we see that when we scale all the linear dimensions equally, the force between permanent magnets goes as the square of length scale.

Since both weight and power-density-limited permanent-magnet linear motor force scale with the cube of distance, this means that the force between permanent magnets becomes relatively stronger as the length scale decreases.

In our experiments with 3.4 cm side-length cubes, parasitic interactions between permanent magnets were very problematic. The permanent magnets would overcome weight and static friction, and lock the system of particles into an invalid configuration with a force too large for the actuators to overcome. Based on the scaling analysis in this section, we could expect these problems to be about 30 times worse at 1 mm scale than 3 cm scale after scaling the system, and thus we are not optimistic about the prospects for millimeter-scale magnetic programmable matter.

2.6.4 Electric-Field Systems

2.6.4.1 The difficulty with Electrostatic “Utility Fog”

In this section, we revisit the “Utility Fog” concept from §2.6.3.1, but using electrostatic rather than magnetic interactions.

The idea is to place capacitor plates on the surface of the nodes, and by applying voltages, move the nodes arbitrarily in three dimensions without contact. Since the nodes do not share a common ground, it is not immediately clear that this is even possible. But even without

considering this question, we can show that the voltages required for certain movements are too high unless the air gap between nodes is small enough.

Consider the voltage required to suspend the cubic particle considered in §2.6.3.1 from a ground plane by applying a voltage between it and an electrode covering the top side of the node. The capacitance is

$$C = \epsilon_0 \frac{D^2}{x} \quad (85)$$

The force between the plates with a constant applied voltage is given by Equation (86).

$$F = \frac{V^2}{2} \frac{dC}{dx} = -\frac{V^2}{2} \frac{\epsilon_0 D^2}{x^2} \quad (86)$$

Combining Equations (85) and (86), making the substitution $x = \lambda D$ to non-dimensionalize the electrode separation, we get Equation (87), which gives the voltage required for suspension.

$$V = \lambda D^{(3/2)} \sqrt{\frac{2\rho g}{\epsilon_0}} \quad (87)$$

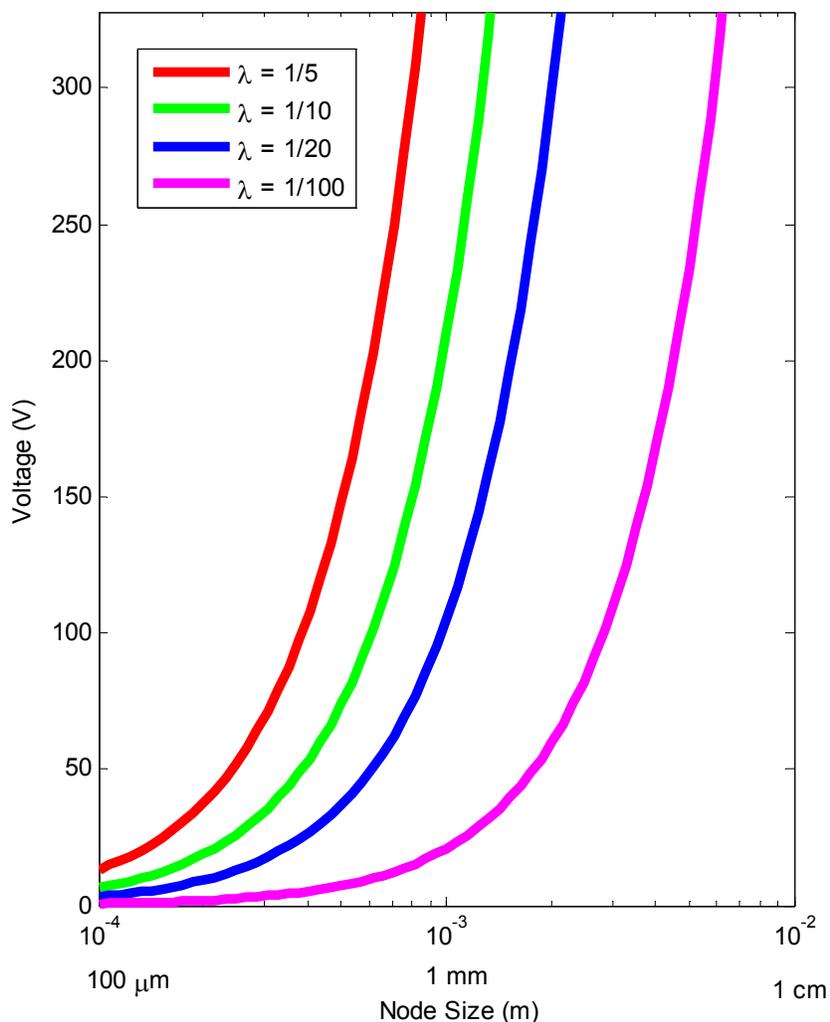


Figure 33: Voltage for Electrostatic Suspension vs. Node Size: This plot shows the voltage required (across a parallel-plate capacitor with a ground-return lead) to suspend a node below a ground plane. λ is the ratio of suspension distance to the node size.

Figure 33 shows the voltage required to suspend a cubic node, with the same assumptions as in §2.6.3.1, using electrostatics. Macroscale electrostatic actuation systems in air are plagued by dielectric breakdown [76], but air cannot break down below 50V at even submicron gaps, due to statistical-mechanical effects. [95] Taking 50V as the maximum safe operating voltage, the separation distance between nodes needs to be less than 5% of the node size for suspension; thus, free-floating “utility fog” is not practical with millimeter-scale particles.

However, if an actuation mode is selected where nodes are always close to contact, with separation distance 1% - 5% of their size, then electrostatic actuation is possible with only modest voltages. Hence, it is our belief that the best hope for millimeter-scale programmable matter lies in close-to-contact electrostatic actuation.

2.7 Integration of Dissimilar Process Technologies

The applications discussed in this thesis require the integration of dissimilar process technologies; for example, the paintable display requires integration of CMOS for logic and III-V processes for visible LED's.

There are several excellent candidate solutions for batch, wafer-scale combination of these processes. They are Wafer-to-wafer bonding [77], III-V fabrication on top of CMOS through SiGe virtual substrates, [78] and epitaxial lift-off. [79]

2.8 Economics

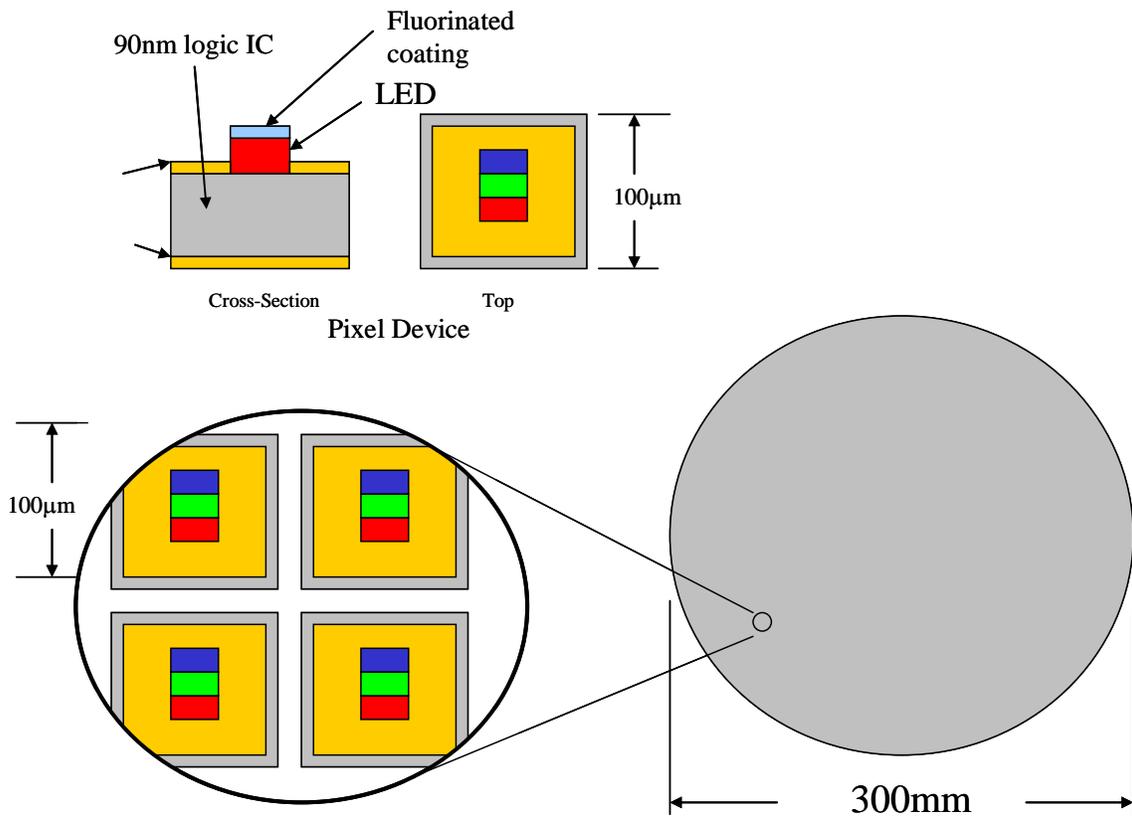


Figure 34: Batch Fabrication of Display Particles

The cost of silicon wafer fabrication, for the leading edge logic technology, has remained approximately constant at \$16/in² for the past 20 years. (1985-2005) [80] This is approximately 2.5 cents per mm². The yield of a semiconductor process goes up with decreasing device size; 1 mm devices can be fabricated with near-unity yield on 300 mm wafers. [81]

The cost of electronic devices sold on the market comes from the cost of packaging, test, assembly, engineering, overhead, and profit; very little of it is due to the wafer fabrication process itself.

Ordinarily, test is a very important step; but programmable matter can tolerate some defective particles. [81]

Here is a generic process flow for the manufacturing of programmable matter particles.

1. Fabricate a silicon wafer with deep-submicron CMOS digital logic.
2. Grow SiGe virtual substrates and fabricate III-V LED's on top of the same wafer.
3. Add any other required materials; magnetic cores for power conversion, for example.
4. Coat the wafer with a protective layer to protect the devices from mechanical damage.
5. Dice (singulate) the wafer by deep reactive ion etching.
6. For a paintable display: mix the resulting millions of dice with a binder and paint onto a surface; for a programmable shape-change system --- pour the dice into a bucket and transmit the power-on command.

The cost of these process steps is given by the cost of the equipment and labor, multiplied by the process time, plus the cost of raw materials.

Fabrication of CMOS wafers is a complicated process involving many steps, highly skilled labor, and very expensive equipment. Thus, it is reasonable to say that the cost of steps 2-7 should not be out of line with the cost of step 1.

For particles without LED's, we estimate the fabrication cost at 1.5 the cost of deep-submicron CMOS fabrication, **3.75 cents / mm²**. For particles with LED's, we estimate the fabrication cost at 2 times the cost of deep-submicron CMOS fabrication, **5 cents / mm²**.

3 APPLICATION EXAMPLES

This thesis is the result of my efforts to design and build working programmable matter particles, for paintable displays and modular robotics. In this section, I describe these two application areas.

For each of the applications, we did a millimeter scale design---not the ones presented here---then built a centimeter-scale prototype capturing as much of the spirit of the millimeter-scale design as possible. The goal of constructing the centimeter-scale prototypes was to learn as much as possible about the chosen designs before attempting microfabrication. Having learned a great deal from building the centimeter-scale prototypes, we reworked the millimeter-scale designs to arrive at the designs presented in this section.

In the end, constraints on time and money did not permit fabricating either of the millimeter-scale designs --- instead this represents exciting future work.

3.1 Paintable Display

A paintable display is an illustrative example of the potential capabilities of passive (actuator-free) programmable matter. A paintable display is sold as cans of paint, and applied with a brush to a surface. The paint contains thousands of semiconductor particles. The particles communicate with one another to determine their relative positions after painting, and work together to render images received from an external data source.

Conventional wisdom is that deep-submicron silicon is too expensive a material with which to construct a large area display --- that is the reason why displays are built using specialized fabrication processes like amorphous-silicon thin-film-transistors (TFT) on glass. However, conventional wisdom assumes that the wafer from which the display is constructed needs to be as large as the display. By making the silicon sparse, with one sub-millimeter chip per pixel, spread out over a large area, with a small area fill fraction, deep-submicron silicon can economically be used to construct a display.

The simplest way to build such a sparse display is to put a light emitter onto each particle. Reflective displays are not compatible with this approach, because reflective displays need a large area fill fraction in order to achieve high contrast. Thus, in our work on paintable displays, we have focused on LED displays.

The centimeter-scale paintable display prototype and distributed rendering algorithms were designed by Bill Butera as part of his PhD thesis work. [13] When I arrived on campus, most of the hardware was already designed and fabricated, and David Greenspan and Monica Sun were already hired and ready to write the software. My contributions to the project were to manage the development of the software, redesign just one of the PCB's, mop up a lot of unexciting hardware loose ends, and chiefly to have the courage to actually turn on the power to this 1000-processor, 80A beast.

3.1.1 Centimeter-Scale Prototype

3.1.1.1 Materials and Methods

The centimeter-scale paintable display prototype contains about 1000 computing nodes, hereafter referred to as “paint particles.” Each particle contains a 32-bit ARM7TDMI processor, 256KB of RAM, 2 MB of FLASH, an infrared optical communications transceiver, a tri-color LED, a light sensor, and a temperature sensor.

The particles are made from four round printed circuit boards arranged in a cylinder, 3.4 cm in diameter, 7.0 cm tall. The particles are placed at random positions on a 4' x 8' foam-core board, and are powered by DC conduction through nails mounted on the underside of the particles, which pierce two metallic power planes embedded in the foam-core board. The communications radius of each particle is intentionally limited to approximately 10 nearest neighbors, using directional optics; an infrared transmitter with a downward-facing conical beam is located on the underside of the second PCB from the top of each particle, and a unidirectional infrared receiver is located on the top side of the second PCB from the bottom of each particle. For more information about the design of the particles, see [81]. The pushpin mechanism described above was designed by Bill Butera in collaboration with Josh Lifton, who constructed similar hardware [82] for a different application.

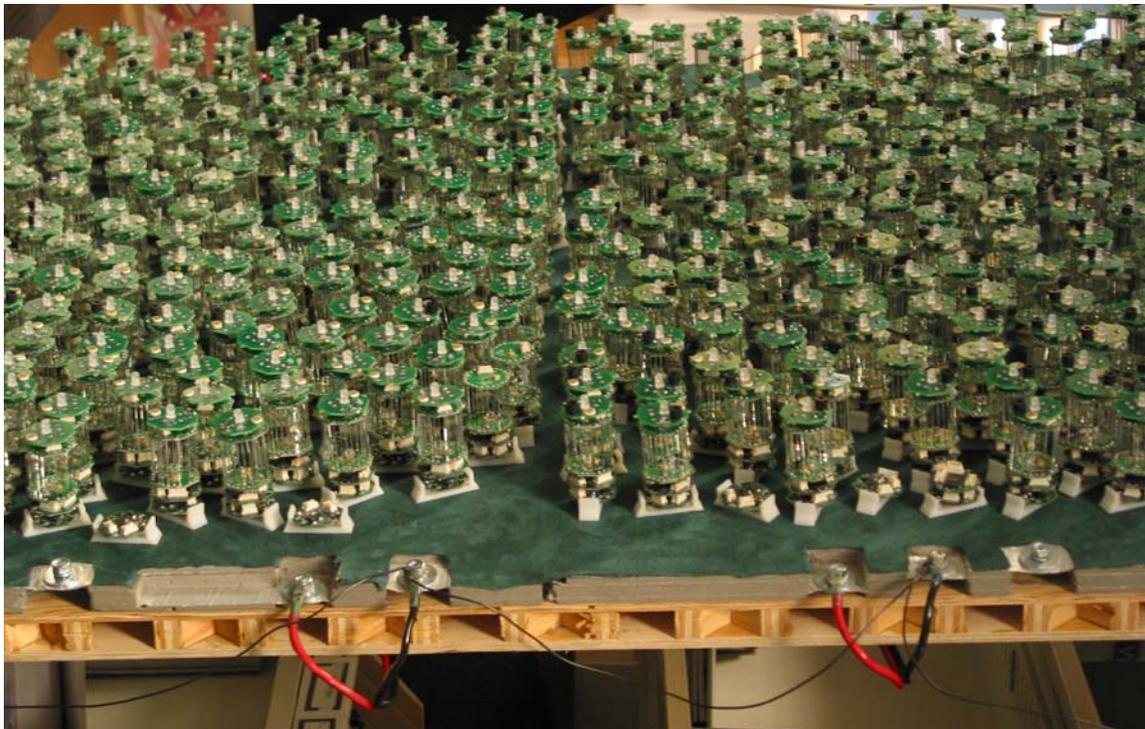


Figure 35: *The centimeter-scale paintable display prototype.*

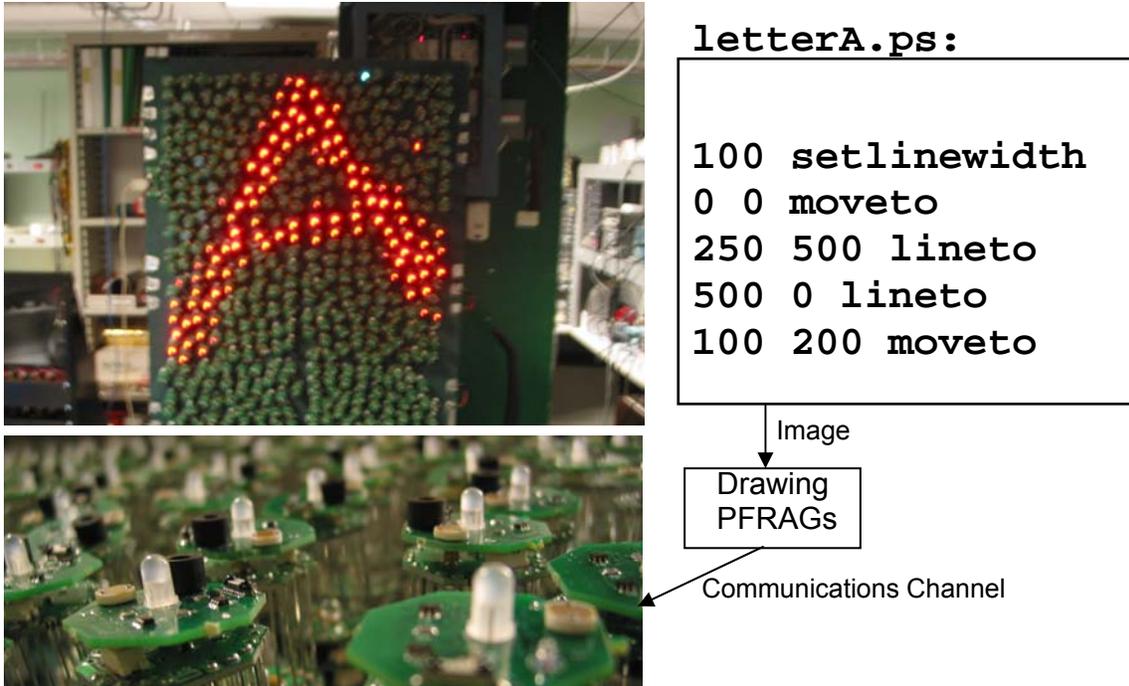


Figure 36: Distributed PostScript rendering on the centimeter-scale paintable display

We built a fully distributed PostScript rendering engine using self-assembling code. [13] PostScript is converted line-by-line into process fragments, which are small binary executables with associated data segments. The process fragments are injected into the system optically. Each particle runs an operating system that executes any properly-formatted process fragment that it receives over its optical communications interface. The process fragments travel from particle to particle to spread through the system. Each process fragment draws one of the objects specified in the PostScript file, by turning on its particle's LED if its particle's coordinates are within half the line-width of the implicit mathematical function defining its object.

At power-up, the particles start out identical, without knowledge of their coordinates. The operator marks three particles as the origin, X axis anchor, and Y axis anchor, by activating their light sensors in sequence with a flashlight. The operator then injects a series of *gradient* process fragments, which allow each particle to approximate its distance to an anchor particle by measuring the minimum communications hop-count from itself to the anchor particle, and then averaging the whole-number hop-counts of neighboring particles to produce a real-valued metric of distance. Finally, the *coordinates* process fragment uses the distance to each of the three anchor particles to triangulate the particle's position in real-valued (x,y) Cartesian space. [81]

3.1.1.2 Results and Conclusion

The performance of the centimeter-scale paintable display was in line with our expectations, and the the algorithms that we tested on our simulator [81] worked well in practice. In the course of building the system, we did discover a few problems with the design, which we summarize here:

A. Defective Particles

About 20% of the particles did not work when we received them from our contract manufacturer. In the paint architecture, occasional bad particles should be acceptable, since the system is self-organizing and assumes random placement and presence of functional particles.

However, for our system, one of the most common particle defects was for the infrared communications LED to be stuck “on.” Particles with defective LED’s flooded the communications channel in their neighborhood, preventing communication between other particles. In addition, a smaller fraction of particles had shorted power supply inputs, pulling down the power supply for the entire board. Both of these classes of defective particles needed to be manually identified and removed, a tedious process.

Because it would be too expensive to 100% test particles in a literal paintable computing system with sub-millimeter paint particles, some defective particles are inevitable. However, we feel that paintable systems with defective particles can be made to work, if the following design principals are followed:

1. Particles should fail fast. Particles should perform a power-on self test. If any defect is detected, the particle should attempt to turn itself off.
2. Particles should be designed with defects in mind. When the particle is designed, a list of the particular ways in which a particle might be defective should be drawn up. The design should be altered so that defects that might affect the system as a whole become less likely.
3. The system should be designed with defects in mind. The communications system and power distribution system should be selected and designed to minimize the chance that a defective particle can interfere with the operation of other particles.

For example, if reactive power distribution had been used instead of conductive power distribution, shorted power supply inputs would not draw down the system bus. If a carrier-modulated infrared system had been used, rather than an on-off keyed system, a stuck-on IR LED would not jam the communications channel. If the circuit traces leading to the IR LED were given more clearance to other traces, the stuck-on LED failure would happen less often. Finally, a fast-acting fuse at the power-supply input terminals would have prevented shorted particles from drawing down the power supply line.

B. Concurrency and Race Conditions

A paintable system is very unforgiving of bad programming practice, such as unchecked race conditions and memory leaks. Because there are 1000 processors running the same code, but with different inputs, any race condition that exists in the code is bound to show up somewhere, in our experience, probably within the first minute after the system is powered.

From our perspective, building a demonstration system, this was a problem. But for an industrial software engineer, aiming to produce 100% reliable code, this would be a feature. Running code on a paintable system exercises it thoroughly in a short time. This makes it easy to reproduce intermittent problems, so that they can be identified and repaired before the software is deployed.

3.1.2 Millimeter-Scale Design

In this section, we present the top-level design of a literally paintable display. This display has the same resolution and area as a 640x480, 17" diagonal, 4:3 aspect ratio computer monitor. But because it is a paintable display, it could also be painted in the form of a 1" x 150" strip, or painted onto the surface of a sphere.

Because it turns out to be too expensive to put a general-purpose processor behind every pixel, we propose to use two kinds of particles; 1 mm particles including a processor and memory, called *rendering particles*, and smaller 110 μm *display particles*, containing light-emitting diodes, a communications receiver, and minimal logic.

3.1.2.1 Display Particles

This display has 307,200 pixels, each responsible for 0.25 mm² of display area. From §2.5.1, 43 μW LED power are required for an indoor-readable display, and 325 μW are required for a sunlight readable display.

Each display particle needs to receive 24 bits of color information from a rendering particle at a rate of 60 Hz. Allowing for 3 times this number of bits, to account for warm-up time, identification bits, and clock skew, display particles need to receive data at 4.3 kb/s.

Because the paint binder needs to contain specialty power-supply materials, it is unlikely to be transparent. Therefore, we select inductive communication over optical communication.

A detailed design example for near-field inductive communication is given in §2.3.2.2.1, between two particles with 1 mm² coils. Here, we consider communication between a 1 mm² rendering particle and a display particle 10 times smaller in linear dimension.

Referring to Equation (30), to keep the same mutual inductance, we need to reduce the communications radius by a factor of 4.7, from 10 mm to 2 mm, to be able to use the results of §2.3.2.2.1.

The display particles can receive data with an efficiency of better than 37 pJ/bit. At the data rate above, this means 0.16 μW per display particle is required for communications.

The digital circuitry in a display particle is shown in Figure 37. We estimate 4 μW for the oscillator and 1 μW for the remainder of the digital logic, which runs at the low speed of 4.3 kHz.

Adding up the power for display, communication, and logic, we get a power budget of **50 μW** for indoor-readable display particles and **336 μW** for sunlight-readable display particles. The heat dissipation limit at 100 μm particle size is about 400 μW; so indoor-readable particles will run cold, and the sunlight-readable particles will run hot.

The 307,200 display particles required consume 3,072 mm² of total silicon area. Using the cost estimate for particles with III-V integration given in §2.8, we can estimate the cost for the display particles at **\$154**.

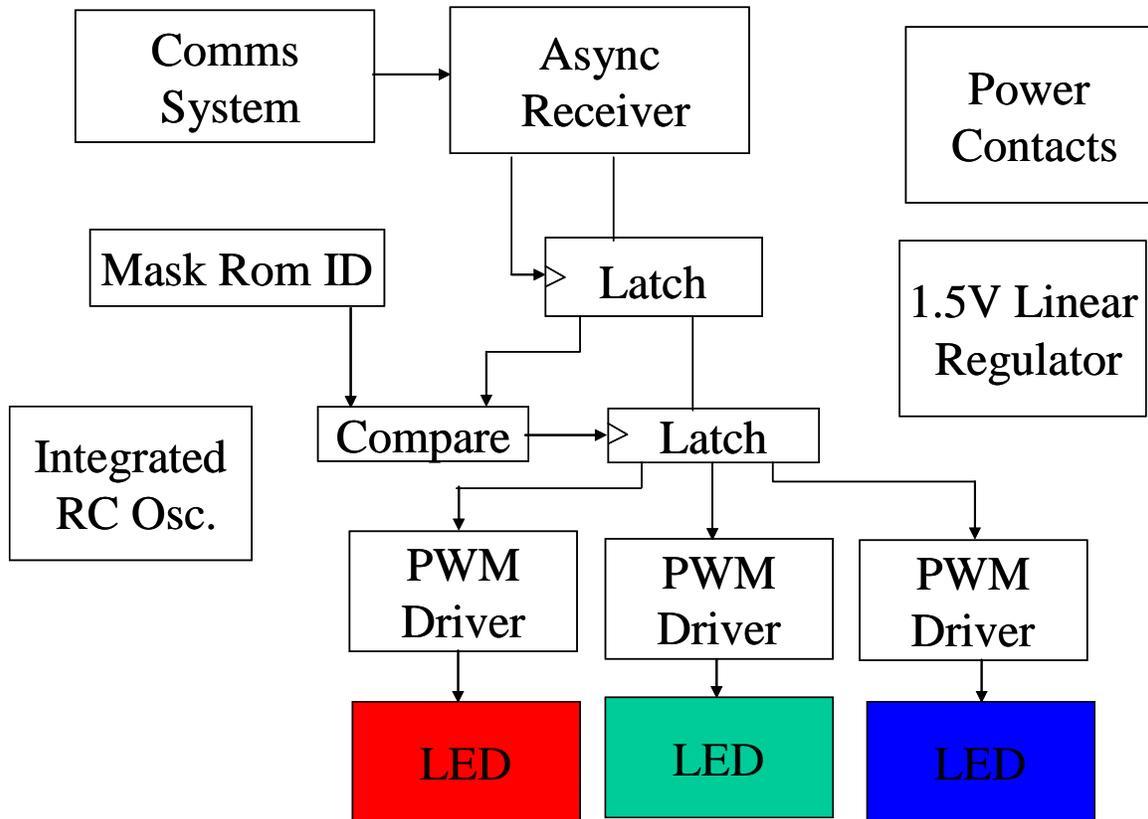


Figure 37: Display Particle Block Diagram: In the proposed paintable display, there are two kinds of particles: 1 mm rendering particles, which have a processor and memory, and 110 - μ m display particles, each of which has a tricolor LED, an inductive communications system, and simple digital decoding circuitry.

3.1.2.2 Rendering Particles

We specify 5,120 rendering particles per display, so that there will be one rendering particle every 3 mm of linear dimension. (With a very uniform coating process, fewer may be required; with a very haphazard coating process, more would be required.) At the density stated, there is one rendering particle for every 60 display particles.

Rendering particles will need to collaboratively decode MPEG streams and PostScript files. The main computational task in MPEG decoding is computing 8x8 inverse discrete cosine transforms, to render 64 pixels of the image. (Each rendering particle is responsible for approximately that number of pixels.) Performing 8x8 IDCT's at 60 Hz takes 35,000 instructions per second. Data transfer and operating system overhead are another computational task; since we cannot accurately estimate this overhead, we will be conservative and guess that processor will need to run at 1 MHz and have 16 KB RAM. Then, rendering particles will draw **60 μ W** each.

The 5,120 rendering particles consume 5,120 mm³ of silicon area. Together, they have 90 MB of static RAM and can execute 10⁹– 10¹¹ operations per second.

Using the cost estimate for paint particles without III-V integration, given in §2.8, we can estimate the manufacturing cost of these particles at **\$192**.

3.1.2.3 Power

Inside each mm² of display area, there are four display particles and 1/9 of a rendering particle. The system power consumption per mm² is

$$50 \mu\text{W} \times 4 \text{ mm}^2 + 60 \mu\text{W} \times 0.11 \text{ mm}^2 = \mathbf{207 \mu\text{W} / \text{mm}^2} \quad (\text{Indoor Readable})$$
$$336 \mu\text{W} \times 4 \text{ mm}^2 + 60 \mu\text{W} \times 0.11 \text{ mm}^2 = \mathbf{1.4 \text{ mW} / \text{mm}^2} \quad (\text{Sunlight Readable})$$

3.1.2.3.1 Power: Random Environment

Using the inside-out zinc-air battery concept detailed in §2.2.2.1, and assuming a 1 mm thick paint layer, 6.0 J/mm² of energy is available from battery reactants stored in the binder. A painted primary battery would allow **eight hours** of indoor operation and **71 minutes** of outdoor operation, for a display of any size.

3.1.2.3.2 Power: Controlled Environment

To power the display from the commercial power grid or from a vehicle electrical system, we can use reactive power transfer, described in §2.2.3.2.

Reactive power transfer can deliver the required power at 8 MHz with $V_{\text{DD}(\text{max})} = 3.3 \text{ V}$, or at 1 MHz with $V_{\text{DD}(\text{max})} = 10 \text{ V}$. A process capable of a least 5 V will be required to fabricate the display in any event, since III-V LED's require that much turn-on voltage.

The oscillating electric field required for reactive power transfer must be produced directly under the display. This could be done by patterning a single-layer flex circuit with interdigitated electrodes at 100 μm pitch, which is readily achievable through low-cost patterning methods.

3.2 Modular Robot

3.2.1 Centimeter Scale Prototype: Magnetic Cubes

In September 2006, we embarked on the design and construction of a sliding-cube shape-change programmable matter system, with centimeter-scale cubes. We chose the sliding cube geometry because it maintains large-area face contact between nodes over an entire reconfiguration, for efficient field-based actuation, as opposed to rotational systems (e.g. spheres, rhombic dodecahedrons) which only have a point or line contact during reconfiguration. For more information about the sliding-cube geometry for modular robots, see [99].

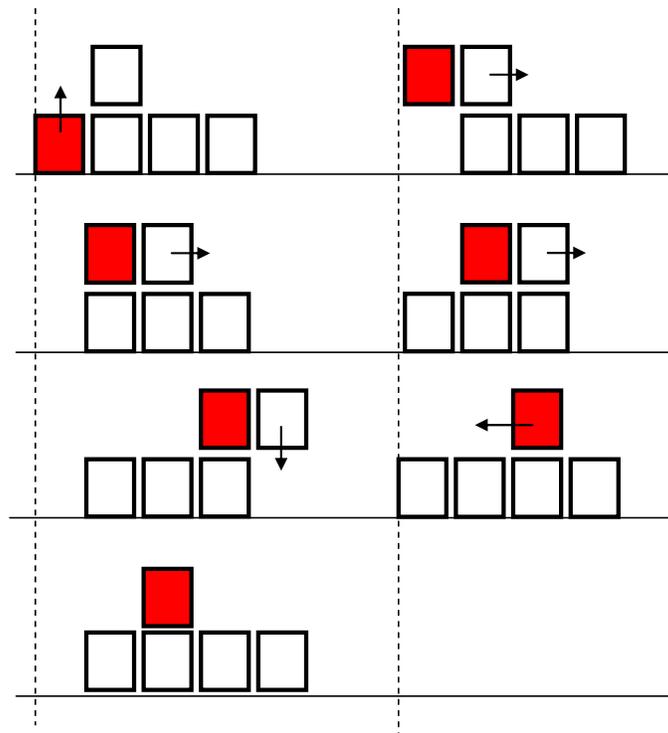


Figure 38: *Sliding cube modular robot locomotion by reconfiguration*

3.2.1.1 The Quandry of Magnet Placement

When we began the design of this system, we wanted to use a conventional magnetic linear motor design. [83] However, we quickly learned that the design constraints on a linear motor used together with a linear bearing are very different from the design constraints on the linear motor faces of a cubic programmable matter node.

The force of attraction between the forcer and stator of a conventional linear motor using iron pole pieces is typically greater than 10 times the peak linear propulsive force. Although the idealized coefficient of friction for Teflon-on-Teflon is 0.04, surface contamination and surface

imperfections act to reduce this in practice, so it is difficult or impossible to run most linear motors without a rolling-element or aerostatic bearing. Rolling elements bearings did not meet our no-moving-parts design goal. Additionally, a typical design constraint for a low-friction bearing is that the length to diameter ratio be greater than two --- otherwise a phenomenon known as “drawer jamming” can occur. [84] This constraint makes it (very) difficult to design a no-moving-parts modular robot structure that keeps an air-gap between nodes, yet still allows for a given node to move arbitrarily through a three-dimensional lattice. As such, we decided to design the nodes without bearings explicit bearings, and attempted instead to design an efficient, scalable, permanent-magnet linear motor with a low attractive force.



Figure 39: A few of the passive magnet-placement prototypes

In order to build a permanent-magnet actuator to move the cubes without power transmission components, the magnet must be on one cube and the current-carrying wires in the other. This requires the placement of permanent magnets near the node surface.

Early in the project, we built passive node prototypes, using laser-cut acrylic and rare-earth magnets, (Amazing Magnets, Irvine CA) to test magnet arrangements for stability. These tests were almost universally disastrous. In general, the problem was that the nodes would not stay where you put them; the permanent magnets gave the system a very sharp lowest-energy state, and the nodes would pull themselves with great force into that configuration. With some configurations, the permanent magnets were far enough apart in valid lattice arrangements that it was possible to align the nodes into a shape and have them hold their positions by static friction, but then the slightest nudge would flip a node 90 degrees, or worse, clump all the nodes into a very un-lattice-like ball.

Using weaker magnets was not a feasible option, because power-efficient actuation depended on strong magnets. What was needed was an arrangement for the magnets in which the desired lattice structure was not too far from the lowest-energy state.

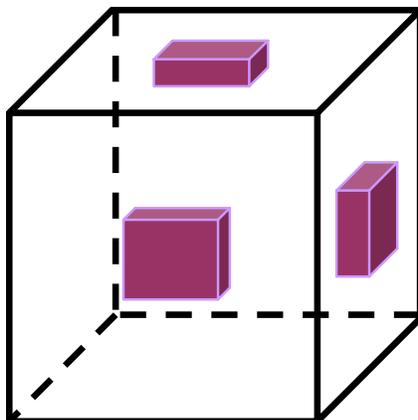


Figure 40: Preferred Permanent Magnet Geometry: Three NIB permanent magnets, magnetized through their thickness, are placed at the center of three orthogonal faces of a cube, with their magnetization direction (North pole) facing out. In this configuration, the cubes naturally align themselves into a geometry close to a cubic lattice, because the North side of one magnet is pulled toward the South side of the corresponding magnet on the next cube.

After some iterative design, we arrived at the arrangement shown in Figure 40. This design works because the north side of all of the magnets faces out, and the south side is protected inside the cube, so very low energy direct North-to-South configurations are geometrically disallowed. It is critically important that the magnets not be too wide or too thick, or else they will interact strongly from the edge as the cubes are slid over one another. At the centers of the faces of a 34 mm cube, 10 mm square magnets, 3 mm thick, with remnant magnetization of 1.2 T, resulted in a reasonably stable cubic lattice, although the result was by no means perfect or ideal; finding a better magnet arrangement would be desirable.

This geometry has the unfortunate aspect of making the nodes asymmetric: nodes have a distinct top, front, and right. Once arranged in a lattice, sliding without rotation from lattice site to lattice site maintains this handedness, but the initial lattice must be carefully arranged by hand.

It might be possible, given proper programming of the nodes, for a group of nodes to be placed in a bag and shaken, and for the actuators to repel nodes that stick onto the lattice with the wrong handedness, so that eventually, the system converges into a valid lattice.

3.2.1.2 Geometric Shielding of Magnets using Nubs

One very good solution to the static magnet placement problem was to shield the rear pole of the magnet using a ferromagnetic cover, and to shield the front of the magnet and cover using geometry, by inseting the magnet at the bottom of a well, protected by protruding nubs. Figure 41 shows this design. We fabricated these units by water-jet cutting (OMAX, Kent, WA) two-part steel magnet housings, assembling rare-earth magnets into the housings, and finally

inserting the assembled magnets into ABS plastic 3-D printed frames. (Stratasys, Eden Prairie, MN) This solution displayed only minor parasitic forces between magnets, much smaller than the weight of the cubes.

However, with this design, the coil needs to sit inside a small, square protrusion to access the magnets (Figure 41, Lower Middle and Lower Right) so it is not in contact with a magnet for the entire reconfiguration motion. In addition, the plastic thickness required for structural integrity, about 1mm, introduces a 1 mm gap between the magnet and coil, vastly decreasing the magnetic field at the coil. Perhaps with additional design and optimization work, and thinner materials, a workable solution might be obtained, but we did not develop this concept further in the interest of time.

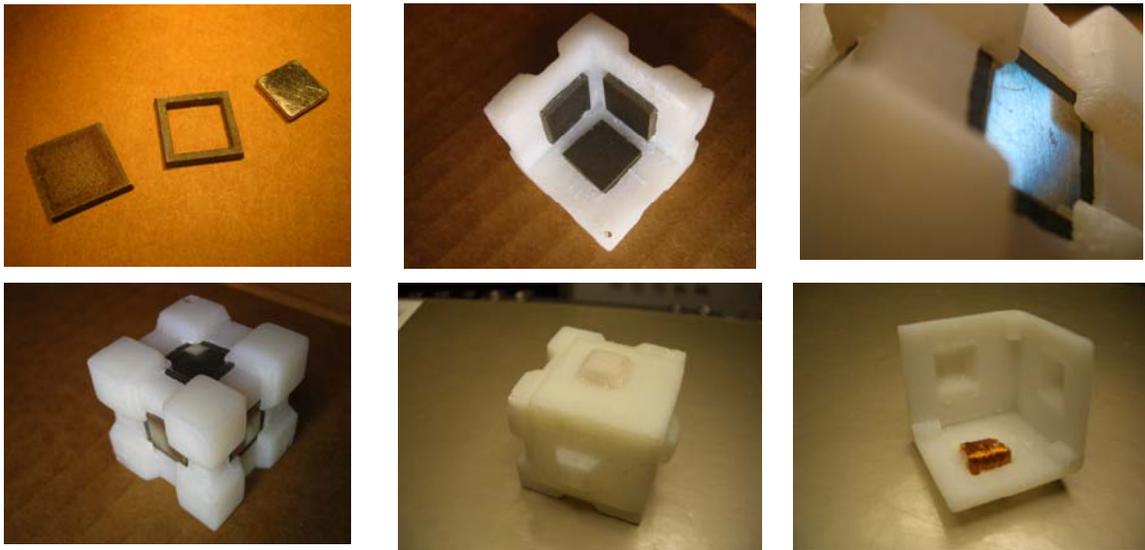


Figure 41: Shielding cubes against permanent magnet interactions using geometry and soft magnetic material: **Upper Left:** A rare-earth permanent magnet (right) and the two pieces of its steel housing. The housing greatly reduces the flux from the rear of the magnet, reducing its interaction with magnets of neighboring cubes. **Upper Middle:** Three magnets installed into one-half of a 3D-printed housing. **Upper Right:** A close up of one permanent-magnet face of a cube showing the installed shielded permanent magnet and the protection nubs, which prevent two magnets from getting too close to one another. **Lower Left:** A complete assembled cube, from the permanent magnet side. **Lower Middle:** A complete, assembled cube from the coil side. **Lower Right:** The space allotted for the coils does not allow efficient actuation.

3.2.1.3 Design of the Coil

The cubes have permanent magnets on three orthogonal faces, and a coil assembly on the three opposing faces.

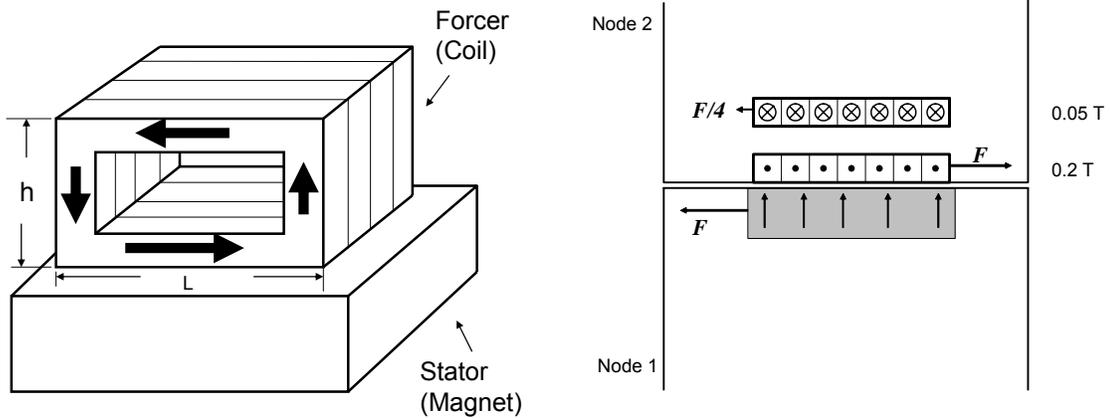


Figure 42: The coil assembly (forcer) of one cube next to the permanent magnet (stator) of another cube.

Left: Perspective view, showing dimensions. **Right:** Cross section view



Figure 43: A photograph of the coil assembly shown schematically in Figure 42.

The coil has four windings (phases), each with bipolar software-controlled drive. The force on a coil segment can be computed from the Lorentz Force law,

$$\vec{F} = \int_V (\vec{J} \times \vec{B}) dV \quad (88)$$

The ohmic power dissipation in the windings can be computed from

$$P = \int_V J^2 \rho dV + \vec{F} \cdot \vec{v} \quad (89)$$

The exact closed-form expression for the magnetic flux density of a rectangular permanent magnet is worked out in Appendix A, allowing computation of the force and moment from this actuator versus displacement, magnet geometry, winding geometry, and coil current.

This actuator produces more-or-less axial force because the magnetic field above the magnet is directed up and the current in the windings is directed laterally. The force is perpendicular to both, so is directed axially. The field at the top layer of windings is much lower than the field at the bottom layer of windings, so there is a net axial force. (See Figure 42)

To optimize the design of the coil to maximize the ratio of force to square-root-power, we ran a multi-parameter, exhaustive search over the parameterized geometry of the coil.

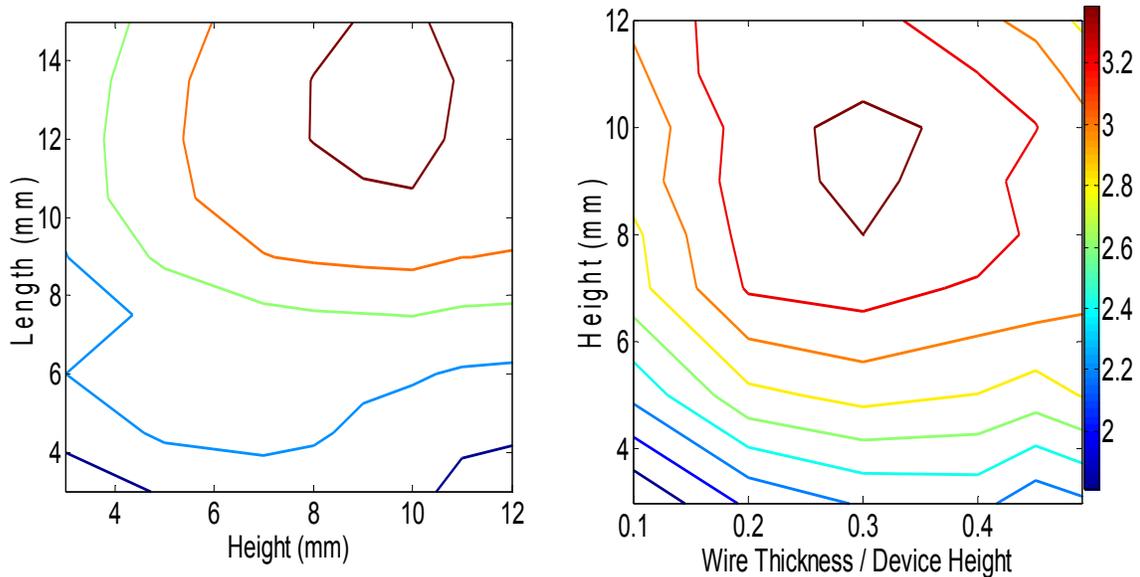


Figure 44: Design charts for a coil section, for 10 mm square NIB magnets, 3 mm thick. **(Left)** The optimum coil is 12 mm long and 10 mm high. **(Right)** The optimum ratio of wire thickness to device height, 0.3, translates to an even distribution of the height of the coil into thirds, with the top and bottom of the coil made of wire, and the middle third made of empty space.

3.2.1.4 Materials and Methods

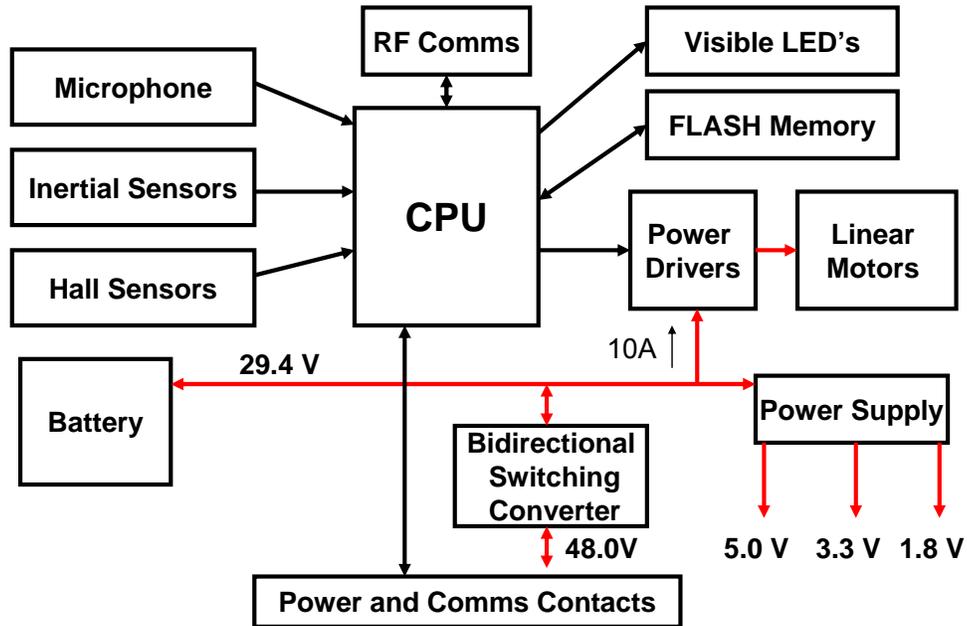


Figure 45: Block Diagram of the node circuitry, described in detail in this section.

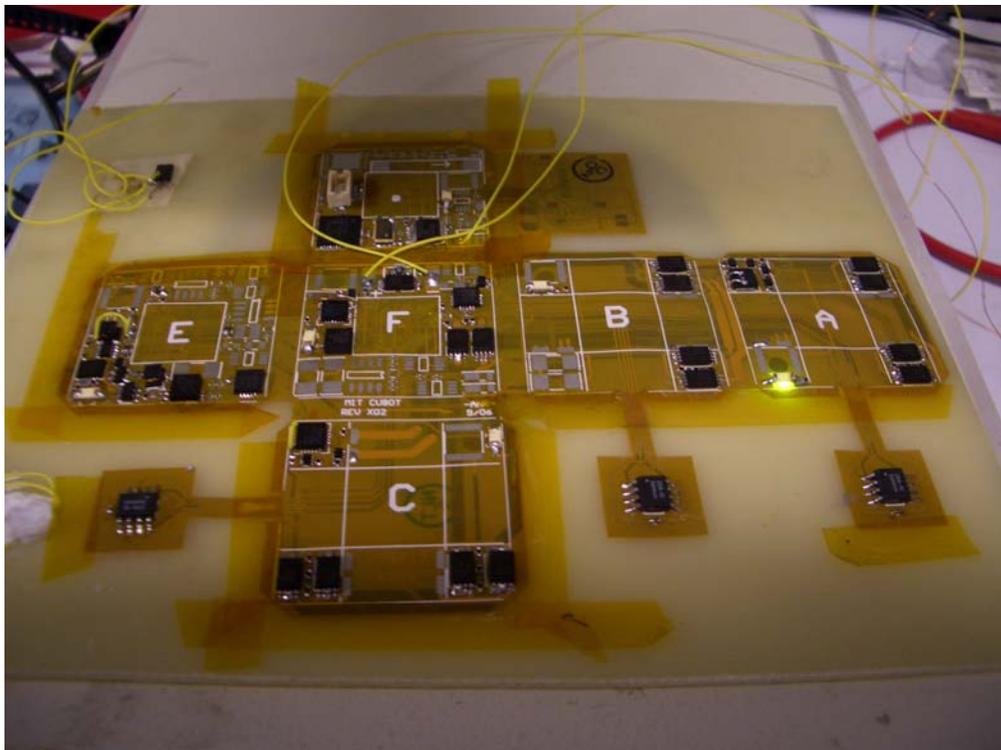


Figure 46: Node Circuit Board being populated: We taped the node circuit boards to an FR4 substrate with Kapton tape for stuffing and initial test.

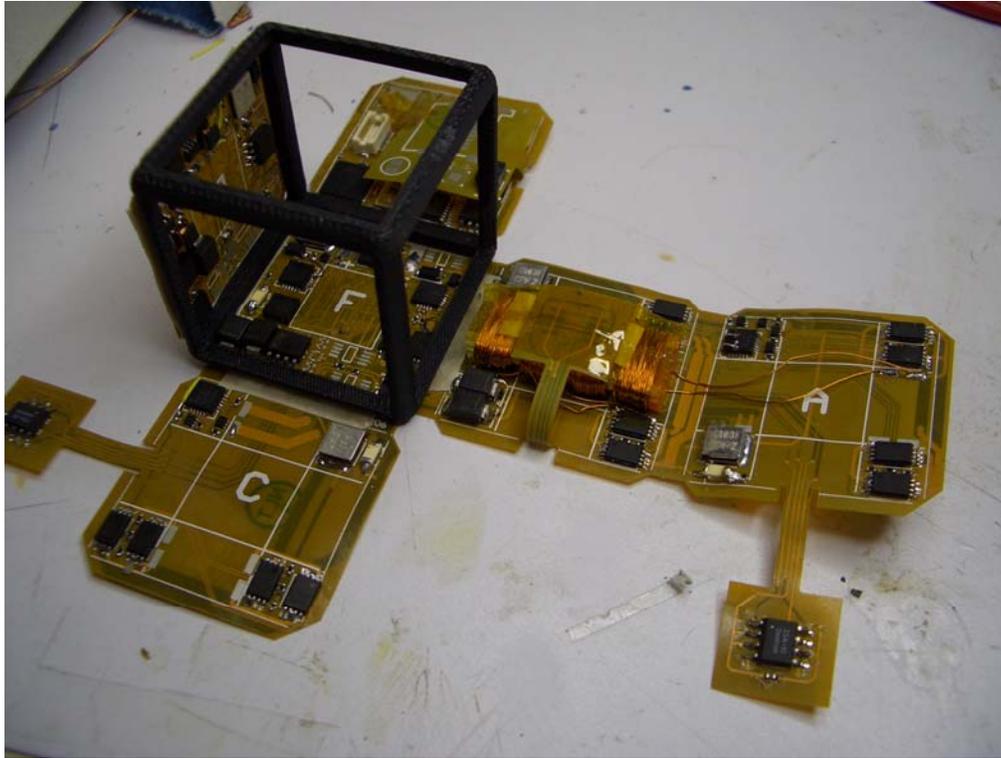


Figure 47: Fully populated node circuit board, partially attached to the frame

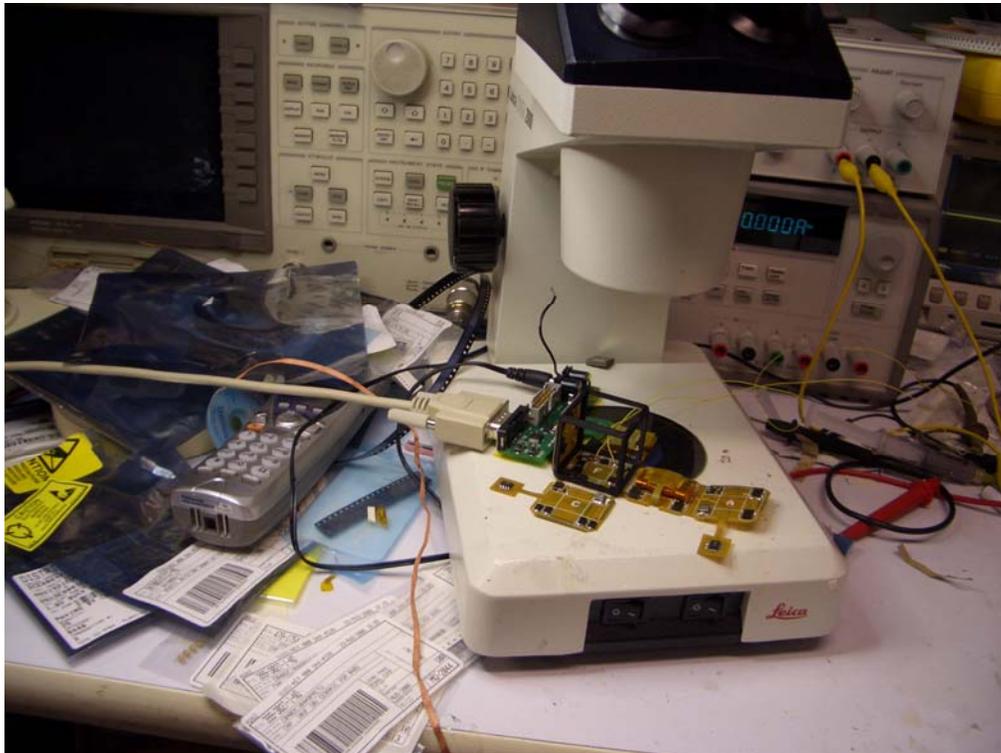


Figure 48: Hardware Development: The node, attached to the PC for programming via a rigid-PCB dongle, on the base of the microscope used for assembly.



Figure 49: A fully assembled node

Each node includes a microprocessor, A/D converter, power supplies, permanent magnets, coil sets, magnetic field sensors, inertial sensors, status indicator lights, and a microphone.

The components used were chosen for their availability in small packages and low power consumption.

The node microprocessor is a LPC2106 (NXP Semiconductor, Eindhoven, The Netherlands), which has an ARM7TDMI core, 128kB flash, and 64kB RAM, and an internal clock rate of 60 MHz. The processor was chosen to be deliberately oversized for the task, to make programming the system maximally simple. The microprocessor is programmed through a JTAG header, connected through a DF12-20 ZIF connector (Hirose, Inc.), which is externally accessible when one of the flaps of the cube is folded back.

Each node has six linear accelerometers. The accelerometers are ADXL320 (Analog Devices Inc., Norwood, MA) dual MEMS accelerometers, with a $\pm 5g$ range, a bandwidth of 2.5 kHz, and a noise density of 250 $\mu g/rtHz$. [85] Each dual accelerometer unit is mounted on a perpendicular face of the flex circuit. Double-integrated acceleration tracks ground-truth position well on the short (100 ms) time scales used for lattice translation.

Each node has six hall-effect sensors. The hall effect sensors are 2SA10 (Sentron AG, Switzerland) dual-axis sensors. A dual-axis hall-effect sensor chip is mounted on a flex-circuit flap above each actuator coil set. The magnetic sensor provides absolute information about a node's relative position, by sensing the fields from the permanent magnets on neighboring nodes.

The six accelerometer readings, six magnetic field sensor readings, the microphone audio, actuator currents, and the power supply voltage are periodically sampled by a MAX1238 A/D converter (Maxim Integrated Products, Inc., Sunnyvale, CA). There are 23 analog signals to convert in total, and the selected A/D converter has only eight inputs, so two ISL43681 8-1 analog multiplexers (Intersil, Milpitas, CA) are used to multiplex lower frequency signals.

Power for each node is supplied by a pack of six 130mAh "Ultra High Discharge Rate" lithium polymer cells. (Atomic Workshop, Norfolk, England) The cells have a mass of 3.2 g each, and have dimensions 25.7 mm x 20.6 mm x 4.67 mm, not including the flexible connection leads. These cells have good high-rate performance --- they can be discharged at 13C (13 times nominal capacity per hour, 1.69 A) and deliver 74% of nominal capacity. In our tests, they are well-modeled as a voltage source in series with an equivalent resistance. The cells have an open-circuit voltage of 4.2 V full charged, and about 3.4 V when depleted, with an approximately linear discharge curve over time between these values. The six-cell pack has dimensions 28.0 mm x 25.7 mm x 20.6 mm. It has an open-circuit voltage from 25.2 – 20.4 V, and a mass of 19.2 g.

The battery power is used directly by the actuator coils, but must be converted to 5V to power the analog section, 3.3V to power the digital section, and 1.8 V to power the microprocessor core. The battery supply is converted to 5V by a LT3470ETS8 (Linear Technology, Milpitas, CA) buck switching converter. The 5V rail is then stepped down to 1.8V by a MAX1921EUT18 buck switching converter (Maxim Integrated Products, Sunnyvale, CA) and the 5V rail is stepped down to 1.8V by a MIC5255 series-pass linear regulator. (Micrel, San Jose, CA)

Power for the actuator coils was switched by complementary MOSFET transistors. The N-channel MOSFET's were IRF7752, the P-channel MOSFET's were IRF7751. (International Rectifier, El Segundo, CA) The N-channel gate was switched from 0-5V. The P-channel gates were driven to battery voltage (off) or half battery voltage (on) by connection to the center of a voltage divider made from 200 ohm resistors, pulled down by a 2N7002 N-channel MOSFET (Fairchild Semiconductor, South Portland, ME) when the P-channel MOSFET is to be turned on, and allowed to float when the P-channel MOSFET is off. The gate of the P-channel MOSFET cannot be pulled all the way to ground because this would violate the datasheet specification on V_{GS} . This circuit was designed for the small packaged size of the components. Although the P-channel gate drive network does dissipate significant power, about 1.5 W, when the P-channel FET is on, this power is dwarfed by the instantaneous power consumed by the actuator coils.

The actuator coils need bipolar drive, but it consumed too much PCB area to assign a full H-bridge to each coil. Instead, each coil is assigned a half-bridge, and the other side is driven by four centralized half bridges that each feed one coil phase from each axis. This prevents more than one axis from being enabled at a time (although they could be time-multiplexed) but cuts the board area for the coil drive subsystem almost in half.

The microprocessor controls the power transistors through three PCA9539BS I²C 16-bit port expanders. (NXP Semiconductors, Eindhoven, The Netherlands)

Each node includes a RF data communications subsystem, centered on a ATA5429, (Atmel Corporation, San Jose, CA) 916 MHz FSK radio transceiver and ANT-916-CHIP chip antenna (Linx Technologies, Merlin, OR)

The coil sets used in the constructed prototype are single-axis. They are fabricated on forms made from laser-cut 1.5 mm-thick acrylic plastic. The forms are 4.8 mm high, and 9.62 mm wide, with each winding pocket 4.5 mm long. The overall width of the coil, including the wire, is 12 mm. The coils are random-wound with a slot-filling quantity of #32 AWG magnet wire around the form, and have a DC resistance of 2.3 Ω .

The coil sets are placed on three mutually orthogonal faces. The other three faces have a rare-earth magnet at the center, 3.175 mm x 9.525 mm x 9.525 mm, magnetized through its thickness, facing north-out.

The magnets and the coils function as linear actuators when the nodes are placed face-to-face. In principle, one could construct a two-dimensional crossed version of the coil by cross-winding the coils at the center, but we did not attempt this for this prototype.

A 29.4 V fully-charged battery pack, discharging 40 ms pulses across one of the 2.3 Ω coils, experiences a voltage drop of 8V, corresponding to 198W of instantaneous power dissipation in the coil.

Each face of the cube has three electrical contacts, laid out at the corners of a square such that when the cubes are placed face-to-face, the contacts touch. Our intention was to place small permanent magnets behind each electrical contact on the flexible circuit, but this proved problematic due to the differing height and geometry of the electrical components on the other side of the flexible circuit. The electrical contacts are POWER, GROUND, and DATA. The data contacts for each side are protected by a 3.3V transient voltage suppressor diode, and each is connected to a microcontroller port pin. The ground contacts are connected to the node's common ground.

The power pin is connected to a power transfer circuit that is designed to allow power transmission through a lattice of connected nodes, so that power provided from an external source can travel through chains of nodes to charge their batteries.

Power from each pin is passed through a SMD030F-2018 positive-temperature-coefficient thermistor (Tyco Electronics, Menlo Park, CA) and SMBJ45A-13-F unidirectional

transient voltage suppression diode (Vishay Electronic GmbH, Germany), which together protect against over-current, over-voltage, and reverse-voltage conditions on the power contacts. This is particularly important for this system, since the contact pads are not mechanically protected.

After the protection stage, power is passed through a 0.15Ω sense resistor, which is monitored by a MAX4081FAUA current sensor (Maxim Integrated Products, Sunnyvale, CA). After sensing, the current is passed through a IRF7342 P-Channel MOSFET (International Rectifier, El Segundo, CA) and connected through to the 48 V power rail.

The 48V rail, which supplies external power, and the 28V battery rail are connected through a bidirectional buck/boost converter. This device is designed to step down from 48V to 28V to draw power from neighboring nodes, or to step up from 28V to 48V to supply neighboring nodes. This design is intended to allow the system to control the amount of power passing from one rail to another, to actively route power through the network.

A slide switch in series with the battery allows nodes to be turned off while not in use. Eventually, we hope that self-charging and an ultra-low-power sleep mode might eliminate the need for a power switch; but we are not there yet in development. A status LED on each face is currently programmed to blink before each move to warn the operator.

3.2.1.5 Results: Volume and Area Usage of Subsystems

The weight of a nodes is an important variable for dynamical simulation, and before constructing a prototype, we could only guess at the weight of the electronics. In the interest of facilitating future theoretical work on this problem, we have included data on the weight, volume, and circuit area usage of our prototype system.

Subsystem	Area or Volume
CPU	A: 58 mm^2
Communications	A: 116 mm^2
Hall Effect Sensors	A: 126 mm^2
Inertial Sensors	A: 196 mm^2
Power Drivers	A: 464 mm^2
Switching Converters	A: 262 mm^2
Actuators	A: 2129 mm^2
Power Transfer System	348 mm^2
Battery	V: 16223 mm^3

Table 9: Area and Volume used by subsystems

The flexible PCB, stuffed with components, has a mass of 6.0 g. Each of the three linear actuator coils has a mass of 3.7 g, and each of the three permanent magnets has a mass of 2.4 g. The battery pack has a mass of 19.2 g, and the cubic frame has a mass of 2.4 g. The total mass of a node is 45.9 g, and the exterior side length is 34 mm, so the overall density of the nodes is **1168 kg/m³**, slightly higher than the density of water. The battery accounts for 42% of the mass, the actuators 40%, and the electronics and frame the remaining 18%.

3.2.1.6 Fabrication Technology

The system was built on a 0.13 mm (5 mil) thick flexible printed circuit with two copper layers and photoimagable coverlay on both sides. (Cirexx International, Santa Clara, CA) The PCB was laid out with 0.13 mm (5 mil) trace/space rules, with 0.56 mm (22 mil) diameter vias. The board had 233 components, and placement and layout took about one week. We assembled the boards using standard solder paste and hot-air reflow techniques. After assembly, we placed magnets and actuators by hand and adhered them to the flex circuit with R9440 double-sided tape (Permacel, East Brunswick, NJ).

We fabricated the cubic frames from ABS plastic using a fused deposition modeling rapid prototyping system. (Stratasys Inc., Eden Prairie, MN) After testing the PCB's for basic functionality, we placed double-sided tape on the bend areas and wrapped them around the cubic frame.

When we assembled the first boards, we had major problems with solder joints breaking around the MLF (micro-lead-frame) components as the board was actively flexed during assembly and use. These problems were solved completely by dispensing a bead of X19331 reworkable edgebond adhesive (Zymet, East Hanover, NJ) around the perimeter of the MLF packages and thermally curing according to the package directions.

3.2.1.7 Experimental Experience with the System

There was generally good visual agreement between the behavior of the system in simulation and in practice, although we made no attempt at precise kinematic measurements from the experimental system.

Nodes are able to successfully translate horizontally. Figure 50 shows a node translating one lattice unit to the right, over a period of 100 ms. The node arrives slightly rotated, and the system has no way to rotate it back, since the actuators have no rotational capability.

Attempts to program a node to translate one lattice unit upward failed, because the actuators exert a torque on the module that rotates it away from the surface. (See Figure 50)

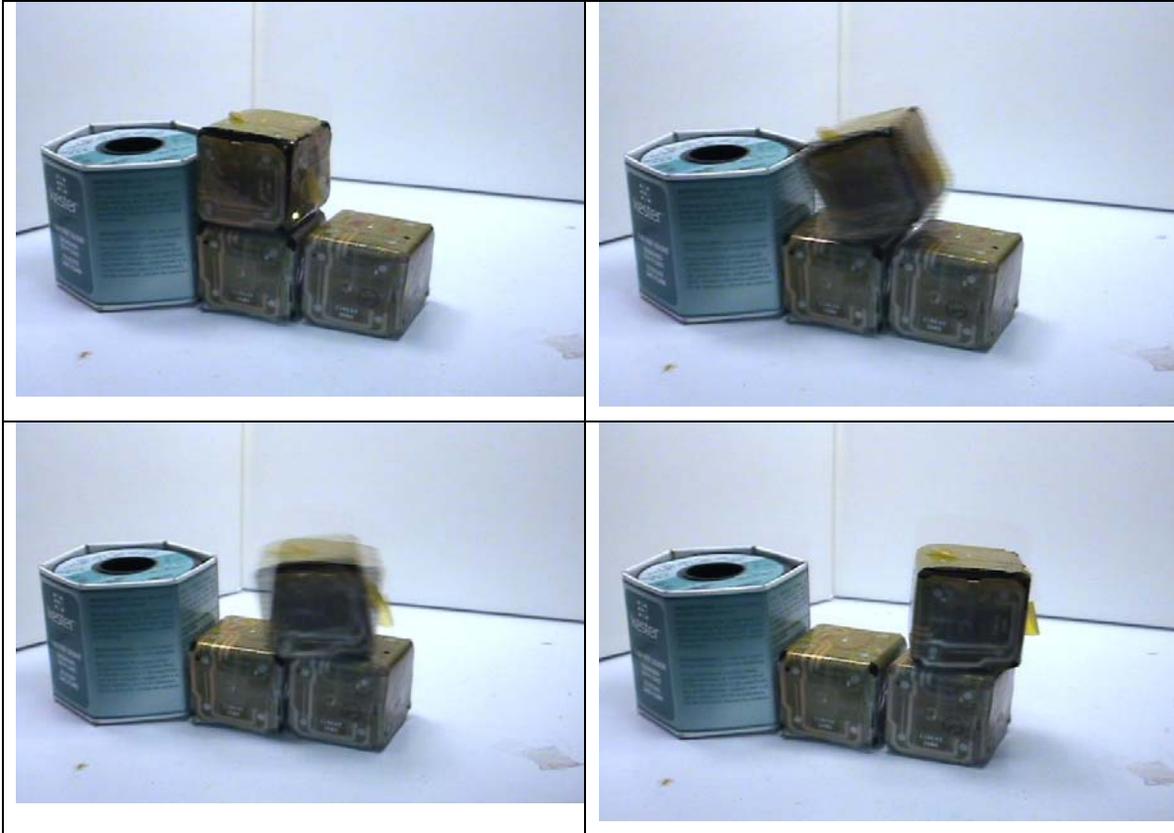


Figure 50: A test of horizontal lattice-unit translation. The images were taken at 33 ms intervals. The top cube contains the coil and power source, the bottom two cubes have permanent magnets at the center of their exposed faces. The top cube propels itself off the first cube, is ballistic for a short time, then pulls itself onto the second cube. The roll of solder is a backstop to keep the bottom cubes from translating to the right during the test.

In the case of horizontal translation, the natural ballistic action of the node after rotation brings it back toward the system for capture by the magnet on the second node. However, for vertical translation up a wall of cubes, the rotational torque causes the node to separate from the wall, preventing actuation.

In the millimeter-scale design described in this thesis, the system has fine-grained control of the force distribution and force direction of the actuators, so this problem could be eliminated with proper control. In addition, it uses electrostatic rather than magnetic actuation, substantially reducing the power consumption of the actuators and weight of the power source, which were significant problem with this design.

3.2.2 Centimeter Scale Prototype: Electrostatic Tiles

We also constructed a centimeter-scale electrostatic system, to learn about the mechanical and electronic behavior of electrostatic surface-drive motors. The electrostatic system consists of tiles that translate on an electrostatic stage. (Figure 51) The entire system was built from printed circuit boards and electronic components.

We implemented three-phase electrostatic motors with interdigitated electrodes, similar to [86]. Motors of this type produce motive force in a direction perpendicular to the electrodes. We fabricated the motors with the electrode pattern shown in Figure 52, to allow the motors to translate in two directions, by switching activation of the electrodes on the stage.

Experimentally, we found that a potential of about 2.6 kV was required to observe electrostatic motor action at all, and 3 kV was required for reliable operation. Because this voltage would cause electrostatic discharge in air, we operated the system submerged under Fluorinert-3283, (3M Corporation, Minneapolis, MN) a dielectric fluid also used for supercomputer cooling.

Figure 53 shows the drive circuitry, which we mounted with standoffs on top of the electrostatic motor board, shown in Figure 52, to form a tile, shown in Figure 51.

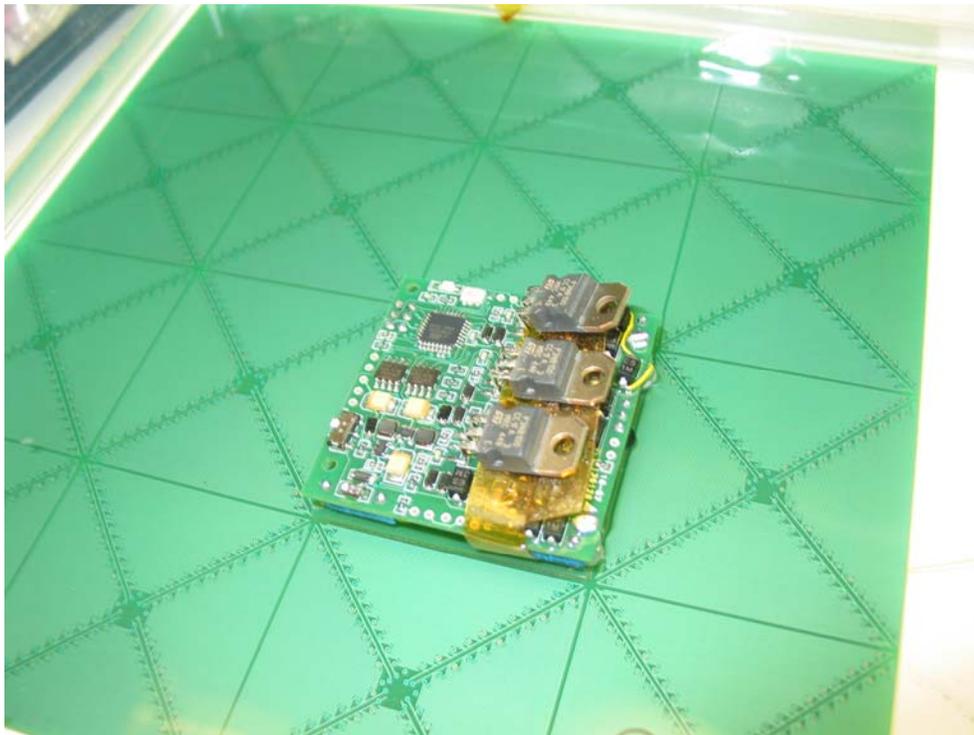


Figure 51: Centimeter-scale electrostatic tile system

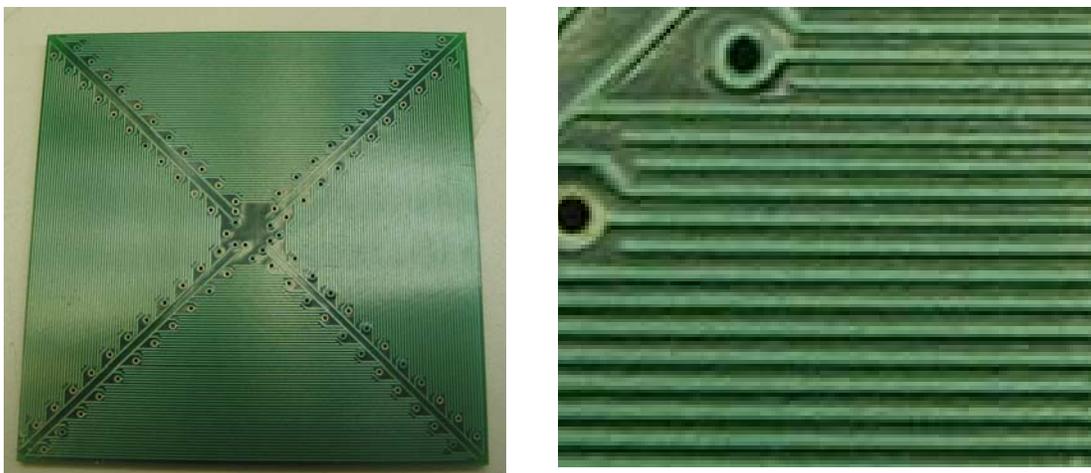


Figure 52: *Electrostatic motor face, overall (left) and close-up (right): The electrodes are placed at a 300 μm pitch.*

Compact, low-cost transistors with V_{DS} rated above 1.2 kV are not available, so it is not a simple matter to build a compact, fast, low-power 3kV switch. Instead, to drive the motors, we built three 3kV transformer-isolated television-flyback power supplies, [87] each driving a 100 M Ω load resistor. The microprocessor energized each of the supplies in sequence to drive the motors, and the resistor bled off the charge once the supply was no longer energized. This arrangement successfully produced the required three-phase set of 3 kV, 10 Hz square waves using about 200 mW of power.

The stage was charged with a static voltage pattern, except that we alternated the polarity of the stage voltage once every 10 seconds, in order to counter what we assume was space-charge buildup on the electrodes from ions in the Fluorinert bath. When we did not alternate the polarity periodically, electrostatic motor action was unreliable and would eventually decay and cease.

Figure 54 is a set of screen captures from a video clip of the node translating. The node moved synchronously with the phase drive, at 2.5 mm/s.

Our experience with this system proved that both electrostatic motor action and static electrostatic latching are possible between separate, electrically floating nodes, an important result for a millimeter-scale design based on electrostatic surface motors.

This system had two major operational problems. The first was that the nodes needed to be rotationally aligned with the stage axes within one electrode pitch for any motion to occur. There was a restoring force that tended to keep the electrodes aligned in the presence of small rotational disturbances, but when large disturbances knocked the electrodes out of alignment, electrostatic/frictional latching occurred and the system had no way to recover.



Figure 53: Driver board, top (**left**) and bottom (**lower**), showing 3 kV flyback power supplies, battery, microcontroller, and support circuitry.

The second problem was that the chosen electrode pattern was not useful for translating more than about a centimeter, about 20% of the size of the electrode pattern. As the node moves, motive force is reduced and static friction is enhanced, because the overlap area between parallel electrodes goes down while the overlap area between perpendicular electrodes goes up. At a certain point along the pattern, motion ceases.

The design of the actuation chips in the millimeter-scale design addresses both of these problems; because it allows for a software-controllable pattern of voltage lines, rather than a fixed pattern, and allows the nodes to accurately determine their relative position and orientation, both of which allow recovery from disturbances.

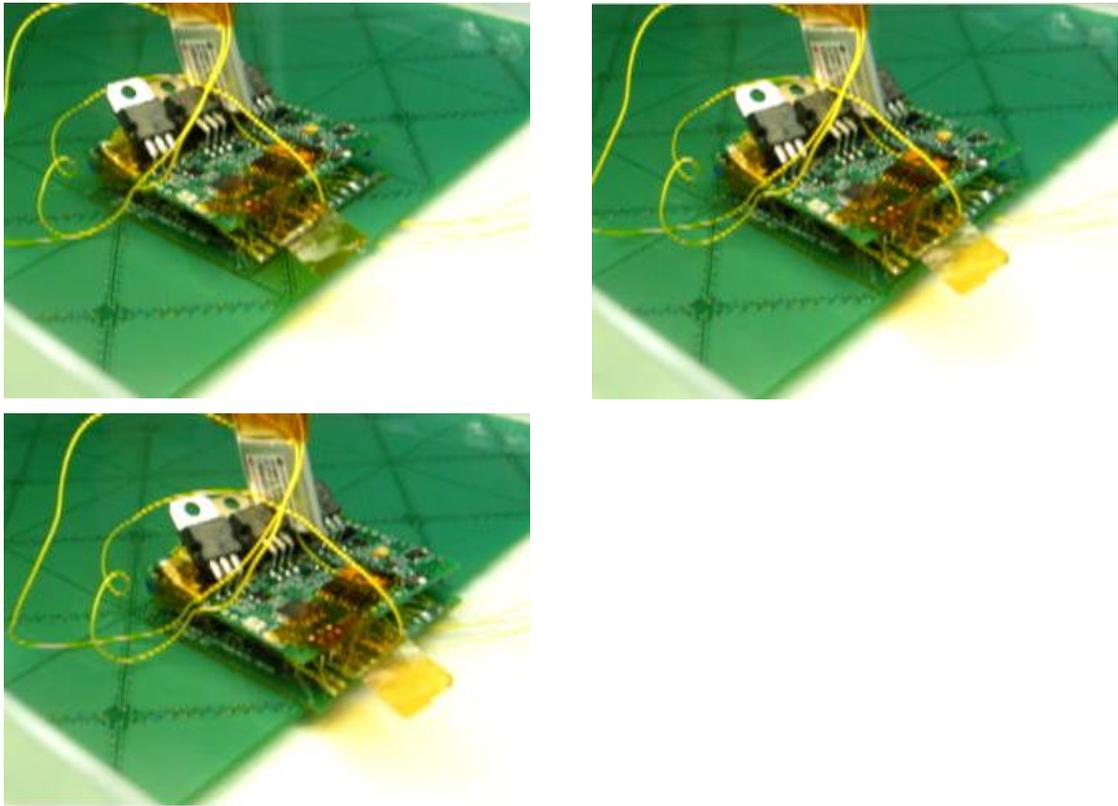


Figure 54: Electrostatic node translating. The motion took about four seconds. The node is electrically isolated from the base plate. (i.e. floating)

3.2.3 Millimeter Scale Design

Figure 55 shows our design for a sliding-cube programmable matter node, with a side length of 2 mm. The node is made from six identical actuation chips, (*Figure 59* and *Figure 60*) one per face, inserted into a space frame. (*Figure 58*) The node also contains a power supply, energy storage capacitor, and microcontroller. Connections between chips are made on the interior of the cube with bond wires. (*Figure 56*)

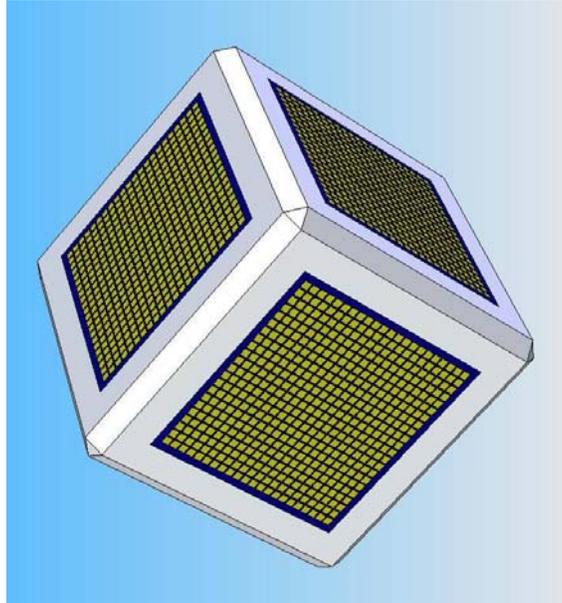


Figure 55: Sliding Cube Programmable Matter Node

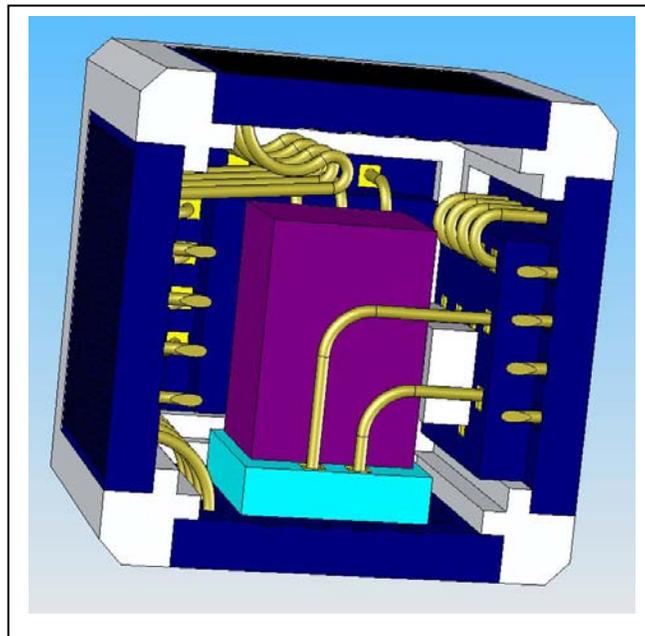


Figure 56: Cutaway view, showing the power source and interconnections between chips.

The actuation chips are designed to operate in face contact with the actuation chips of other nodes. They provide electrostatic actuation, electrostatic latching, capacitive power transfer, capacitive communication, and capacitive localization. The microcontroller controls the operation of the node at a high level.

The power supply charges the + 40V storage capacitor with incoming power from neighboring nodes on an AC bus, and supplies regulated +3.3V and +1.2V for the low-voltage electronics. The fundamental reason for the multiple power rails is a desire for high fields at the actuation surfaces to maximize motive force, but low fields in the computational core to minimize switching losses.

3.2.3.1 Energy Storage

Energy for the node is stored in a standard 0603, 1 μ F, 50V ceramic capacitor. (587-1257-1-ND, Digi-Key Corporation, Thief River Falls, MN) The capacitor is attached to the +40V rail used for actuation power. If the capacitor is charged to 40 V and allowed to discharge to 35 V, the useful energy stored by the capacitor is

$$E = \frac{1}{2} C (V_2^2 - V_1^2) = 0.19mJ \quad (90)$$

3.2.3.2 The Power Supply

The power supply chip is an integrated charge pump that transfers power from the AC bus to the +40V storage capacitor, and from the +40V bus to the +3.3V and +1.2V busses, with linear post-regulation. The charge pump operates at 10 MHz and uses integrated MIM capacitors. It has off-the-shelf 01005 discrete capacitors mounted directly on the die for decoupling; these are visible in Figure 56.

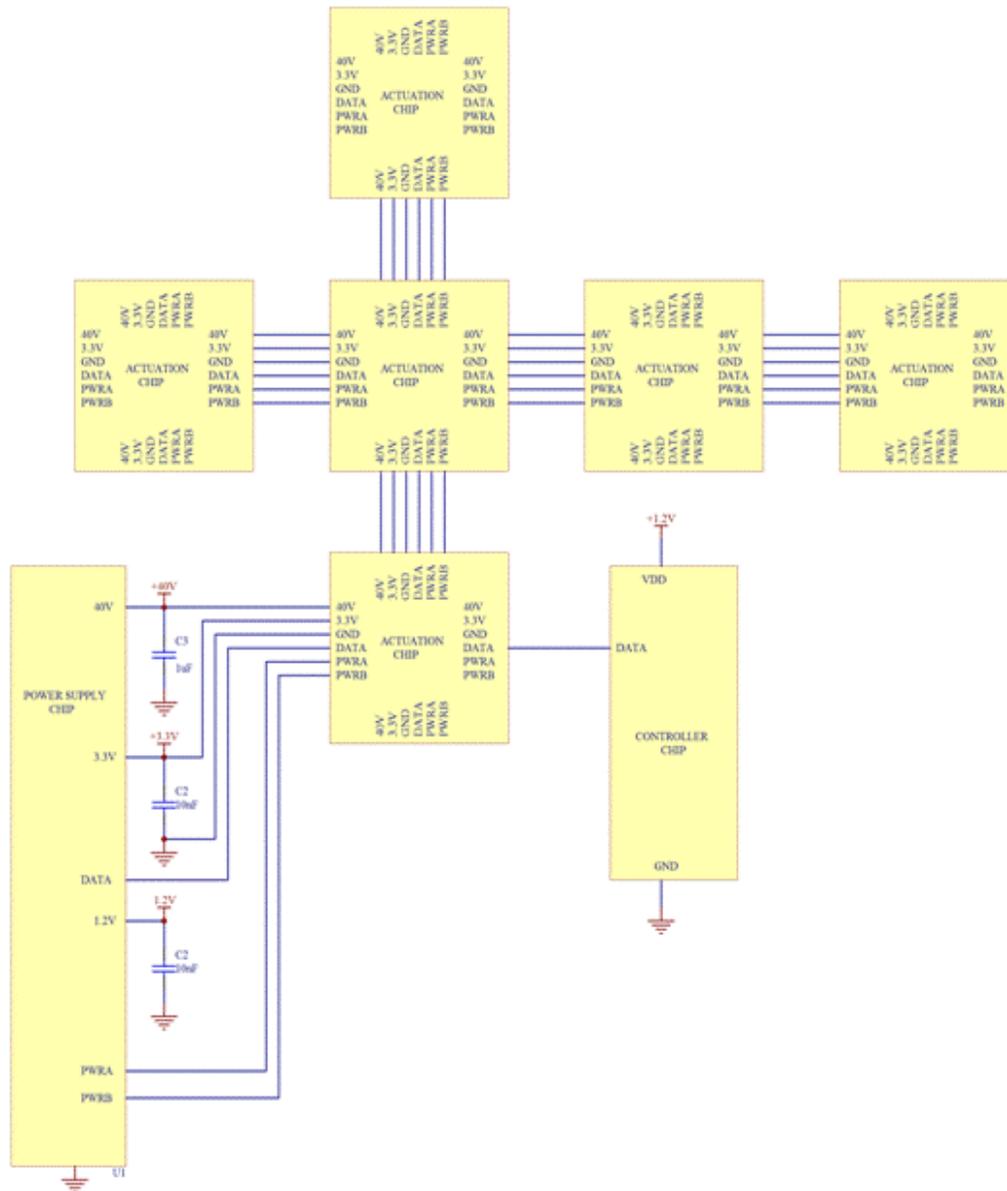


Figure 57: Top-Level Node Schematic showing wirebond interconnections between chips. Each actuation chip determines its address on the one-wire data bus by its relative position in the chain.

3.2.3.3 The Frame

The frame (Figure 58) gives three-dimensional structure to the node. The edges are chamfered to allow nodes to slide over one another without catching. Each face has a mounting position for an actuation chip, 1.4 mm on a side and 200 μ m thick. The frame can be fabricated using the EFAB multi-layer selective electro-deposition process. (Microfabrica Inc., Van Nuys, CA)

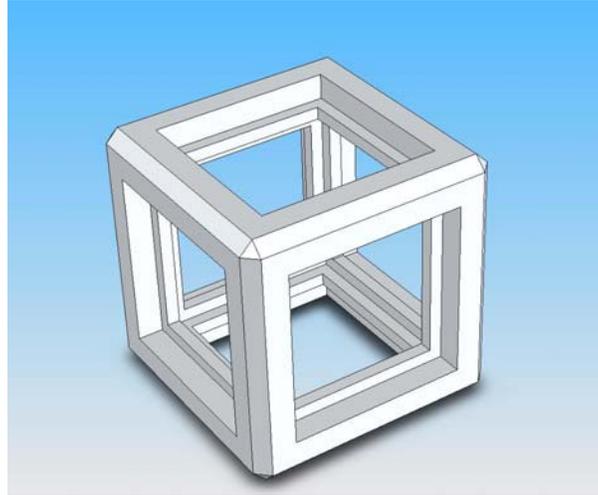


Figure 58: Space Frame

3.2.3.4 The Actuation Chips

The actuation chips (Figure 55, Figure 59 and Figure 60) are fabricated on the Jazz Semiconductor (Newport Beach, CA) 180 nm HV-CMOS 40V display driver process. (Table 10)

Parameter	Value
Features	6 metal, 3 thick metal
Feature size	180 nm
Masks	30
Maximum V_{ds}	40V
MIM Capacitor	2+2 fF/ μ m ²
$R_{ds(on)}$ HV-NMOS (Ω - mm ²)	0.35
LV-MOS Operating Voltage	3.3

Table 10: CA18HR HV-CMOS process specifications

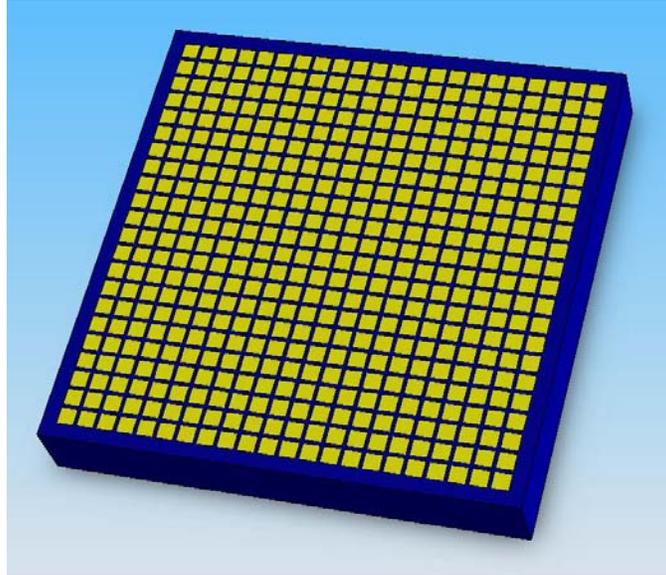


Figure 59: Actuation Chip, Front View, showing an array of actuation pads. A reduced number of electrodes are shown for visual clarity – the design calls for a 110 x 110 array of electrodes.

The actuation chips are responsible power transfer, communication, localization and actuation. Each chip has a 110 x 110 array of actuation cells, placed on a 12 μm pitch. Mechanically, each actuation cell consists of an square electrode on the top metal layer, and its drive circuitry on the lower five metal layers. There is a continuous, unbroken layer of Silicon Nitride covering the top of the chip; the electrodes are not exposed.

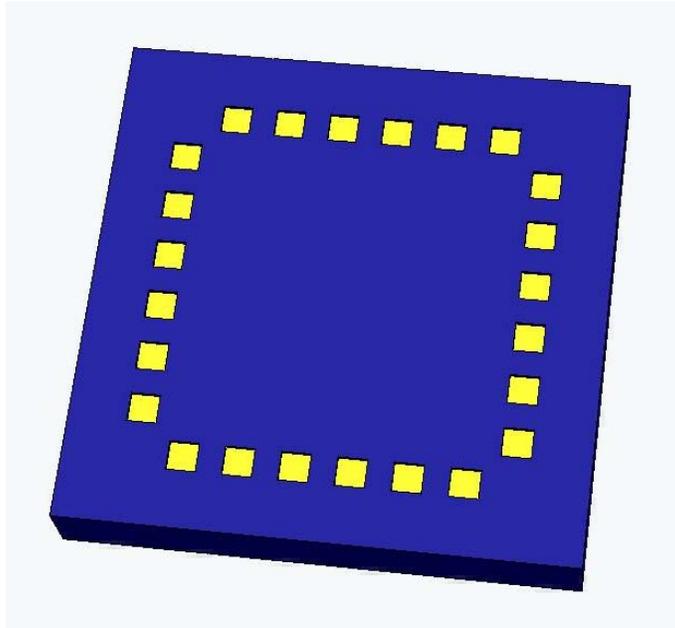


Figure 60: Actuation Chip, back view, showing the backside bonding pads for interconnection with other chips in the same node.

A circuit diagram of an actuation cell is shown in Figure 62, and a block diagram of the entire chip is shown in Figure 61.

Each actuation cell is essentially a CMOS microcontroller port pin driver. Software can set each electrode to 40 V or ground, disconnect the output driver so the electrode acts as an input, or connect the output to one of two AC power distribution busses. Cells at the center of the chip are connected to the PWRA bus, and cells around the edge of the chip are connected to the PWRB bus.

Actuation chips are designed to operate face-to-face. When adjacent electrodes are both driven as outputs, mechanical actuation and/or latching result. When one chip's electrodes are driven and the other are sampled this allows communication and localization via electrostatic fiducials.

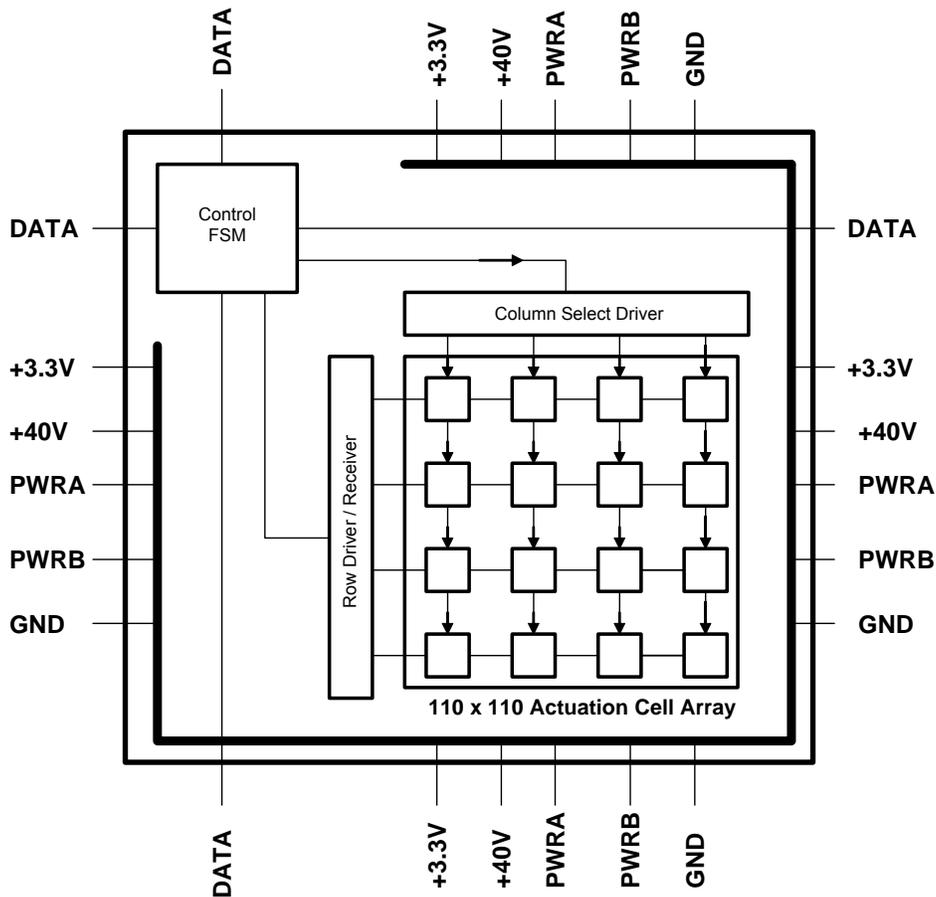


Figure 61: Actuation Chip, Block Diagram

The actuation chips have four sets of power and data pads for connection to other actuation chips in the node. Power is bussed through the chip, and communication is routed

between chips on a one-wire bus by an on-chip finite-state machine. The actuation cells are addressed with a column decoder and a row driver/receiver shift register, in similar fashion to SRAM or a display backplane.

A 180 nm process has a typical standard cell place-and-route density of 110,000 gates/mm², so especially with hand-routing, the 14 gate actuation cell circuit should easily fit inside the 12µm square area of a cell.

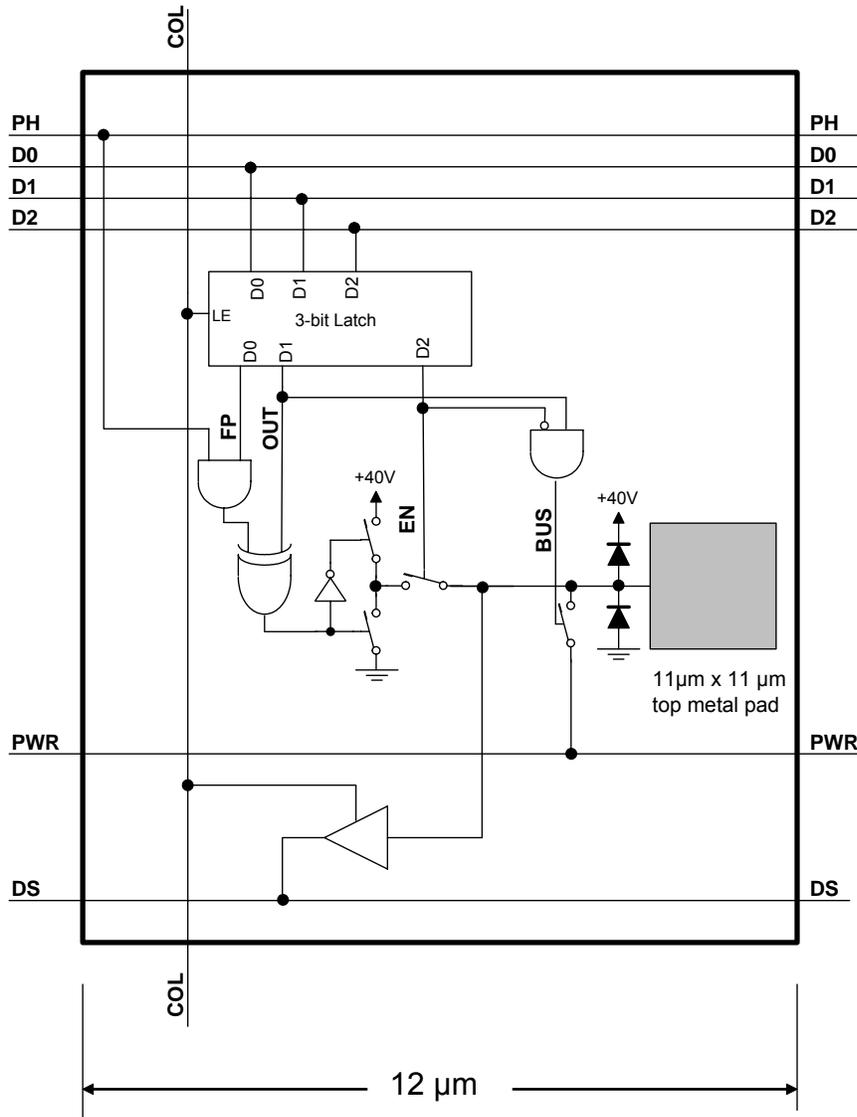


Figure 62: Actuation Cell Block Diagram. Each Actuation Chip Contains a 110 x 110 square array of these cells.

After the actuation chip dice are returned from the foundry, they need to be thinned to 200 μm , and have 24 via holes, 50 μm in diameter, opened to the first metal layer from the back to accommodate wirebonds to power and data lines from the rear of the chip. Wafer thinning is typically carried out using mechanical grinding. [88, 89] Wafer thinning is simplest when carried out on a whole wafer, but by attaching dice to a carrier wafer, individual dice (as would be received from multi-project service) can also be thinned. [90] After wafer thinning, via holes to the first metal layer can be opened by laser ablation of the die [91] or deep reactive ion etching. [92]

3.2.3.5 Latching Force Calculation

200 mm diameter silicon wafers are typically flat to within 15 μm , so over the area of a 1.4 mm actuation chip, we can expect 0.1 μm flatness. The Silicon Nitride passivation layer above the electrodes is 1 μm thick. So the gap between electrodes of two facing chips contains 2 μm of Silicon Nitride, relative dielectric constant 7.5, and 0.1 μm air.

The force per unit area on parallel plates separated by a solid dielectric, in air, is given by

$$\frac{F_n}{A} = - \frac{\epsilon_0 V^2}{\left[(d-b) + \frac{\epsilon_0}{\epsilon} b \right]^2} \quad (\text{from [93]}) \quad (91)$$

where d is the distance between the plates, b , is the thickness of the dielectric, ϵ is the permittivity of the solid dielectric, ϵ_0 is the permittivity of free space, and V is the applied voltage.

Using this equation, the normal force per unit area is 95 kPa. This is the tensile strength of a programmable matter object; it is in the same range as polyimide foam and weak engineering ceramics.

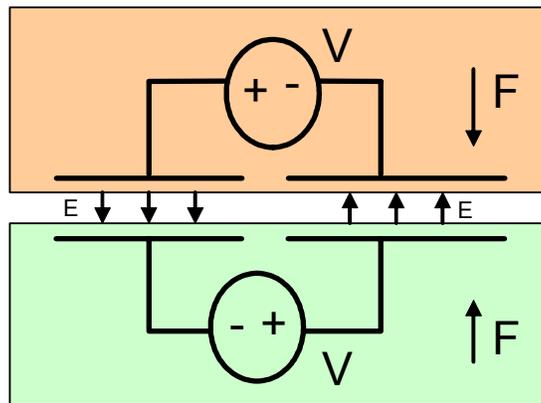


Figure 63: Equivalent circuit for electrostatic latching between two nodes.

The voltage across the Silicon Nitride layer limited to 40V, and 500 V are required for breakdown across a 1 μ m layer. [94] The voltage across the air gap, using the capacitor division law, should average about 10 V during latching, but could be higher in transient conditions. A potential of 40 V across the air gap would be close to, but not necessarily above, the breakdown limit of air at 0.1 μ m. [95]

3.2.3.6 Motive Force Calculation

When motive (shear) force is desired the electrodes are configured into lines to operate as an electrostatic motor in the desired direction, in the manner of the three-phase electrostatic surface-drive film actuators described in [96].

If the electrode pitch is too large relative the gap, the static friction due to the attractive force between electrodes overwhelms the motive force and motion is not possible. As the electrode pitch becomes smaller, fringing field interactions become more important. At a small enough electrode pitch, the direction of the normal force reverses and causes electrostatic levitation, allowing smooth electrostatic motor action. [96] (See Figure 64)

$$\frac{F_s}{A} = \frac{f_s}{f_{s0}} \frac{\epsilon V^2}{p^2} \quad \text{from [96]} \quad (92)$$

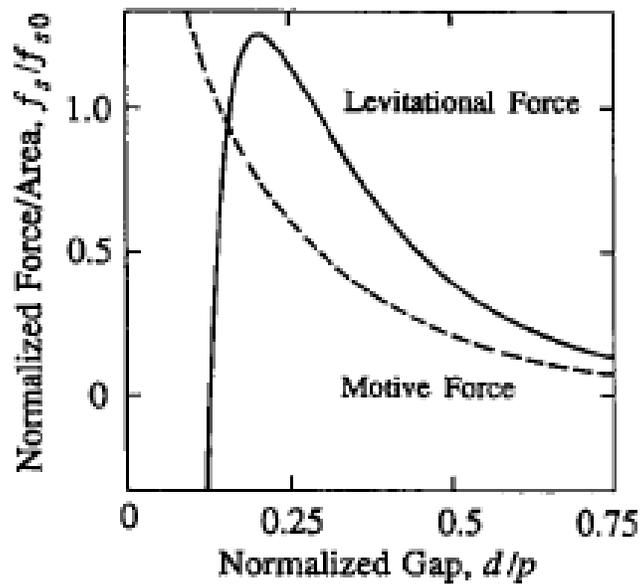


Figure 64: Electrostatic motor force, from [96]. These results are for a three-phase electrostatic induction motor, but should also apply to the motors described here due to very similar principle of operation.

If we select $d/p = 0.16$, since the gap distance between electrodes d is $2 \mu\text{m}$, then the required pitch p is $12 \mu\text{m}$. In this case, using the electrostatic motor scaling equation above, the motive force per area is 550 Pa .

A node has an active area of 1.96 mm^2 , so it can experience a motive force due to the actuators of 1.1 mN , and can exert a latching force of 186 mN . A node weighs 30 milligrams , (16 mg frame, 8 mg electronics, 6 mg capacitor) so gravitation exerts a force of 0.3 mN . Thus, the linear actuators can exert about four times the force needed to lift a single node, and a latching force needed to hold up a vertical chain of 620 nodes.

3.2.3.7 Shaping the Force Distribution

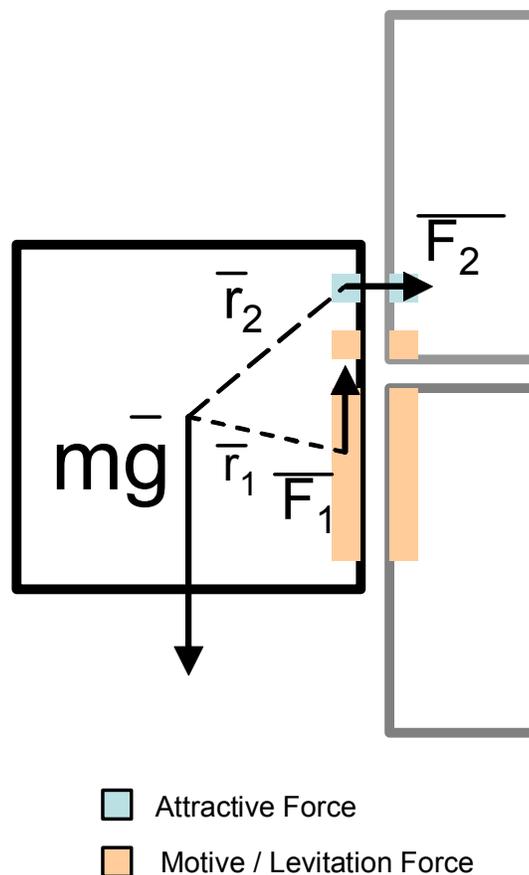


Figure 65: *Mechanics of a cube free-climbing a wall of cubes. Because the actuation chips allow shaping of the force distribution arbitrarily under software control, an attractive force F_2 can be applied to keep the cube against the wall. The magnitude of F_2 would be sized that the torque $r_2 \times F_2 = -r_1 \times F_1$.*

The actuation chips allow the force distribution to be shaped to achieve a variety of motions, by spatially varying the voltage pattern. For example, consider the problem of a single node free-climbing a wall of nodes.

The weight of the node acts through its center of mass, so has no moment, but the motive force F_1 has an approximately perpendicular moment arm, tending to flip the cube counterclockwise off the wall. However, because the actuation chip can have a selectable combination of latching and actuation, it can be programmed to exert a latching force F_2 near the top of the moving cube, which will cancel the moment from F_1 and allow the cube to climb the wall.

3.2.3.8 Motion Speed and Power Consumption

Nodes move 12 μm per step. If the actuation electrode matrix is updated at 100 Hz, nodes can travel at 1.2 mm/s, so they can move the 2 mm to a new lattice position in 1.7 seconds.

The capacitance of a capacitor of area A , made from two dielectrics with thickness d_0 and d_1 is:

$$C = \frac{\epsilon_0 \epsilon_1 A}{\epsilon_0 d_1 + \epsilon_1 d_0} \quad (93)$$

Using this equation, the total capacitance at the interface between two nodes is then 45 pF. The worst-case switching loss at the interface (which can only be charged every other cycle) is given by:

$$P = \frac{1}{4} CV^2 f \quad (94)$$

From this computation, the switching loss at the electrode interface during actuation is 1.8 μW . The mechanical power required to lift a node against its weight is:

$$P = Fv = mgv \quad (95)$$

Carrying out this computation for our 30 milligram node being lifted at 1.2 mm/sec, the mechanical power required to lift a node is 0.4 μW .

To perform one step, a 110 x 110 matrix of electrode cells needs to be updated with 3 bits of data. So the data rate to a single actuation chip is 3.4 Mb/s.

We can expect that low-voltage power consumption on the actuation chip will be dominated the power required to charge each of the 36,300 data latches 110 times for each cycle of the 100 Hz actuation clock. The switching energy for transistors on a 180 nm process is about 110 fJ, [97] so the power consumed to drive the data latches is about 40 μW . Interestingly, this is dramatically higher than the power to charge the electrode capacitances themselves.

An ARM7TDMI 32-bit microprocessor consumes 60 μW / MHz on a 90 nm process, at 1.2 V, and 0.10 mm^2 of silicon area. [98] For headroom, we would like to be able to run the processor 10 times faster than the rate needed to feed the actuation chip with data. Since it is a 32-bit processor and executes roughly one instruction per clock, this requires a clock rate of 1.1 MHz, and 66 μW power.

The sum of the power demands in this section is 110 μW . Since 0.19 mJ is available from the capacitor, this happens to be exactly 1.7 seconds of power, enough to move from one lattice position to the next.

Most of the power dissipation in this system comes from the microprocessor, and from loading data into the actuation array. Very little is actually used for actuation. Custom-designed computational hardware (e.g. a finite state machine or cellular automaton) to run the actuation array might allow the node processor to be powered down during reconfiguration, saving considerable power.

3.2.3.9 Localization

The node processor can configure electrodes on the actuation chips as inputs and read out their state. Thus, it is possible to detect electrostatic fiducials displayed by a neighboring node to recover the relative position and orientation of the two nodes. By propagating position and location information through a network of nodes via electrostatic data communication, each node can determine its orientation and position relative to every other node.

3.2.3.10 Reconfiguration Geometry and Motion Planning

The algorithm needed to transform a sliding-cube modular robot between geometries has been worked out by Butler, Kotay, and Rus in [99], and tested in simulation. The algorithm computes the motions required using a cellular automaton distributed over the nodes.

3.2.3.11 Capacitive Power Transfer

The power flow between two nodes through the actuation chips is limited by:

$$P = \frac{1}{2} CV\Delta Vf \tag{96}$$

Where V is the transfer voltage, ΔV is the allowable ripple voltage at the receiver, and f is the switching frequency. Each face has 1.98 mm^2 of electrode surface, so dividing a face into two capacitors for complete-circuit AC power transfer, each of the two capacitors can have 0.99 mm^2 of area. The capacitance of each is 23 pF. At 40 V, with an output ripple of 5 V, and switching at 10 MHz, this allows the transfer of up to 22 mW, much more than required.

4 **CONCLUSIONS**

Power and communications do present limits on what kinds of applications can be realized using a programmable matter. Nonetheless, many interesting applications are feasible.

Paintable displays appear to be feasible. Battery-powered, field-paintable displays would have a battery life in the 8-hour range. Factory-coated displays could use external power. Both variants could be made at a cost competitive with medium-size LCD's on an area basis, but could be scaled to any size. In addition, they could conform to unusually shaped 3-D surfaces, and could flex.

This thesis presents a design for millimeter-scale shape-change programmable matter, with no moving parts in the nodes, amenable to construction with standard IC foundry process, with minimal post-processing. The design uses a programmable electrostatic matrix actuator, made with a standard high voltage CMOS IC process, to allow face-to-face actuation in any direction, communication, and localization, and power transfer. The design appears feasible from considerations of fabrication time and cost, electrical power, mechanics, heat dissipation, required computational capacity, and required silicon area.

When we started this project, we assumed, based on intuition from work at the macro-scale, that micro-robotics would be a high power application for autonomous microsystems, more so than sensing, displays, or radio communication, due to the requirement for mechanical actuation, and therefore might not be feasible due to limitations on power consumption or heat dissipation. In fact, electrostatic actuation is one of the lowest-power functions possible for an autonomous microsystem, using less power than computation by an order of magnitude.

5 APPENDICIES

5.1 The Magnetic Field of a Rectangular Permanent Magnet

A rare earth permanent magnet has the following B/H relationship:

$$\vec{B} = \mu_0 \vec{H} + \vec{B}_r \quad (97)$$

Where B_r is the remnant flux of the magnet, a vector aligned with the direction of magnetization and with a constant magnet depending on the material, typically about 1.2 Tesla for NIB.

To compute the field and flux density in the region outside the magnet, we can model the magnet as a set of infinitely many magnetic dipoles, aligned with the magnet's magnetization direction, so that the top face of the magnet has a uniformly distributed sheet of positive magnetic charge, and the bottom face has a uniformly distributed sheet of negative magnetic charge.

The magnetic flux vector at a point due to a magnetic charge is

$$\vec{B}(\vec{r}) = \mu_0 \frac{q_{m1}}{|\vec{r}|^3} \vec{r} \quad (98)$$

The magnetic flux vector at a point due to a sheet of magnetic charge is

$$\vec{B}(\vec{r}) = \int_A \mu_0 \frac{\sigma_m}{|\vec{r} - \vec{r}_q|^3} [\vec{r} - \vec{r}_q] dA \quad (99)$$

In Cartesian coordinates, for a square sheet of magnetic charge centered at a height z' above the origin in the X/Y plane, with width w in the X direction and length l in the Y direction:

$$\vec{B}(x, y, z) = \int_{-w/2}^{w/2} \int_{-L/2}^{L/2} \mu_0 \frac{\sigma_m}{[(x-x')^2 + (y-y')^2 + (z-z')^2]^{3/2}} \left[\hat{i}(x-x') + \hat{j}(y-y') + k(z-z') \right] dy' dx' \quad (100)$$

First working with the y-direction term, we have

$$\vec{B}_y(x, y, z) = \mu_0 \sigma_m \int_{-w/2}^{w/2} \int_{-L/2}^{L/2} \mu_0 \frac{(y-y')}{[(x-x')^2 + (y-y')^2 + (z-z')^2]^{3/2}} dy' dx' \quad (101)$$

We can integrate this in y using the indefinite integral

$$\int \frac{a-y'}{[(a-y')^2 + b]^{3/2}} dy' = \frac{1}{[(a-y')^2 + b]^{1/2}} \quad (102)$$

Resulting in

$$\vec{B}_y(x, y, z) = \sigma_m \mu_0 \int_{-w/2}^{w/2} \left[\frac{1}{[(x-x')^2 + (y-y')^2 + (z-z')^2]^{1/2}} \right]_{-L/2}^{L/2} dx' \quad (103)$$

Evaluating this, we get

$$B_y(x, y, z) = \sigma_m \mu_0 \int_{-w/2}^{w/2} \frac{1}{\left[(x-x')^2 + (y-L/2)^2 + (z-z')^2 \right]^{1/2}} - \frac{1}{\left[(x-x')^2 + (y+L/2)^2 + (z-z')^2 \right]^{1/2}} dx' \quad (104)$$

We can do this integral in x' using the indefinite integral

$$\int \frac{1}{\left[(a-x')^2 + b \right]^{1/2}} dx' = -\log(a-x + \sqrt{(a-x)^2 + b}) \quad (105)$$

Resulting in

$$B_y(x, y, z) = \sigma_m \mu_0 \left[-\log\left(x-x' + \left[(x-x')^2 + (y-L/2)^2 + (z-z')^2 \right]^{1/2} \right) + \log\left(x-x' + \left[(x-x')^2 + (y+L/2)^2 + (z-z')^2 \right]^{1/2} \right) \right]_{-w/2}^{w/2} \quad (106)$$

Substituting

$$B_y(x, y, z) = \sigma_m \mu_0 \left[-\log\left(x-L/2 + \left[(x-W/2)^2 + (y-L/2)^2 + (z-z')^2 \right]^{1/2} \right) + \log\left(x-L/2 + \left[(x-W/2)^2 + (y+L/2)^2 + (z-z')^2 \right]^{1/2} \right) + \log\left(x+L/2 + \left[(x+W/2)^2 + (y-L/2)^2 + (z-z')^2 \right]^{1/2} \right) - \log\left(x+L/2 + \left[(x+W/2)^2 + (y+L/2)^2 + (z-z')^2 \right]^{1/2} \right) \right] \quad (107)$$

Taking all the terms inside a single log, we can write

$$B_y(x, y, z) = \sigma_m \mu_0 \log \left[\frac{\left(x-W/2 + \left[(x-W/2)^2 + (y+L/2)^2 + (z-z')^2 \right]^{1/2} \right) \left(x+W/2 + \left[(x+W/2)^2 + (y-L/2)^2 + (z-z')^2 \right]^{1/2} \right)}{\left(x-W/2 + \left[(x-W/2)^2 + (y-L/2)^2 + (z-z')^2 \right]^{1/2} \right) \left(x+W/2 + \left[(x+W/2)^2 + (y+L/2)^2 + (z-z')^2 \right]^{1/2} \right)} \right] \quad (108)$$

Exchanging the order of integration in equation 4, we see that the equation for B_x is identical to that for B_y , so long as one swaps the names of the x and y variables, and the names of the parameters W and L . So we can write down the expression for B_x by inspection from the expression for B_y .

$$B_x(x, y, z) = \sigma_m \mu_0 \log \left[\frac{\left(y-L/2 + \left[(y-L/2)^2 + (x+W/2)^2 + (z-z')^2 \right]^{1/2} \right) \left(y+L/2 + \left[(y+L/2)^2 + (x-W/2)^2 + (z-z')^2 \right]^{1/2} \right)}{\left(y-L/2 + \left[(y-L/2)^2 + (x-L/2)^2 + (z-z')^2 \right]^{1/2} \right) \left(y+L/2 + \left[(y+L/2)^2 + (x+W/2)^2 + (z-z')^2 \right]^{1/2} \right)} \right] \quad (109)$$

Now, considering the z direction, which is physically different from the x and y directions because it is the magnetization direction, and mathematically different because we do not integrate with respect to dz' , we have

$$\bar{B}_z(x, y, z) = \mu_0 \sigma_m \int_{-w/2}^{w/2} \int_{-L/2}^{L/2} \mu_0 \frac{(z-z')}{\left[(x-x')^2 + (y-y')^2 + (z-z')^2 \right]^{3/2}} dy' dx' \quad (110)$$

We can integrate this with respect to y' using

$$\int \frac{1}{\left[(a-y')^2 + b \right]^{3/2}} dy' = -\frac{a-y'}{b \left[(a-y')^2 + b \right]^{1/2}} \quad (111)$$

Resulting in

$$\bar{B}_z(x, y, z) = \mu_0 \sigma_m \int_{-w/2}^{w/2} \left[\frac{-(z-z')(y-y')}{\left[(x-x')^2 + (z-z')^2 \right] \cdot \left[(x-x')^2 + (y-y')^2 + (z-z')^2 \right]^{1/2}} \right]_{-L/2}^{L/2} dx' \quad (112)$$

Evaluating, we get

$$\bar{B}_z(x, y, z) = \mu_0 \sigma_m (z - z') \int_{-w/2}^{w/2} \left[\frac{(y - L/2)}{\left[(x - x')^2 + (z - z')^2 \right] \cdot \left[(x - x')^2 + (y - L/2)^2 + (z - z')^2 \right]^{1/2}} + 0 \frac{(y + L/2)}{\left[(x - x')^2 + (z - z')^2 \right] \cdot \left[(x - x')^2 + (y + L/2)^2 + (z - z')^2 \right]^{1/2}} \right] dx' \quad (113)$$

We can integrate this expression using

$$\int \frac{1}{\left[(a - x')^2 + b \right] \cdot \left[(a - x')^2 + c \right]^{1/2}} dx' = - \frac{\tan^{-1} \left((a - x') \sqrt{\frac{c - b}{b \left[(a - x')^2 + c \right]}} \right)}{\sqrt{b(c - b)}} \quad (114)$$

Yielding

$$\bar{B}_z(x, y, z) = \mu_0 \sigma_m (z - z') \left[\frac{(y - L/2) \tan^{-1} \left[(x - x') \sqrt{\frac{(y - L/2)^2}{(z - z')^2 \left[(x - x')^2 + (y - L/2)^2 + (z - z')^2 \right]}} \right]}{\sqrt{(z - z')^2 (y - L/2)^2}} + \frac{(y + L/2) \tan^{-1} \left[(x - x') \sqrt{\frac{(y + L/2)^2}{(z - z')^2 \left[(x - x')^2 + (y + L/2)^2 + (z - z')^2 \right]}} \right]}{\sqrt{(z - z')^2 (y + L/2)^2}} \right]_{-w/2}^{+w/2} \quad (115)$$

Simplifying

$$\bar{B}_z(x, y, z) = \mu_0 \sigma_m \operatorname{sgn}(z - z') \left[\frac{+ \operatorname{sgn}(y - L/2) \tan^{-1} \left[\frac{(x - x') |y - L/2|}{|z - z'| \sqrt{(x - x')^2 + (y - L/2)^2 + (z - z')^2}} \right]}{- \operatorname{sgn}(y + L/2) \tan^{-1} \left[\frac{(x - x') |y + L/2|}{|z - z'| \sqrt{(x - x')^2 + (y + L/2)^2 + (z - z')^2}} \right]} \right]_{-w/2}^{+w/2} \quad (116)$$

Where we have made use of the sgn function to simplify (a-b)/sqrt((a-b)^2). Now, evaluating this expression, we get

$$\bar{B}_z(x, y, z) = \mu_0 \sigma_m \operatorname{sgn}(z - z') \left[\frac{+ \operatorname{sgn}(y - L/2) \tan^{-1} \left[\frac{(x - W/2) |y - L/2|}{|z - z'| \sqrt{(x - W/2)^2 + (y - L/2)^2 + (z - z')^2}} \right]}{- \operatorname{sgn}(y + L/2) \tan^{-1} \left[\frac{(x - W/2) |y + L/2|}{|z - z'| \sqrt{(x - W/2)^2 + (y + L/2)^2 + (z - z')^2}} \right]} - \frac{- \operatorname{sgn}(y - L/2) \tan^{-1} \left[\frac{(x + W/2) |y - L/2|}{|z - z'| \sqrt{(x + W/2)^2 + (y - L/2)^2 + (z - z')^2}} \right]}{+ \operatorname{sgn}(y + L/2) \tan^{-1} \left[\frac{(x + W/2) |y + L/2|}{|z - z'| \sqrt{(x + W/2)^2 + (y + L/2)^2 + (z - z')^2}} \right]} \right] \quad (117)$$

Because of the equivalence between x and y in the problem set-up, this expression must have the property that the names of x and y and W and L can be swapped and result in the same

expression. In fact, this is the case --- because arctangent is an odd function, we can push the sgn expression into the arctangent to achieve the symmetrical in x/y expression for z:

$$\bar{B}_z(x, y, z) = \mu_0 \sigma_m \left[\begin{array}{l} + \tan^{-1} \left[\frac{(x-W/2)(y-L/2)}{(z-z')\sqrt{(x-W/2)^2 + (y-L/2)^2 + (z-z')^2}} \right] \\ - \tan^{-1} \left[\frac{(x-W/2)(y+L/2)}{(z-z')\sqrt{(x-W/2)^2 + (y+L/2)^2 + (z-z')^2}} \right] \\ - \tan^{-1} \left[\frac{(x+W/2)(y-L/2)}{(z-z')\sqrt{(x+W/2)^2 + (y-L/2)^2 + (z-z')^2}} \right] \\ + \tan^{-1} \left[\frac{(x+W/2)(y+L/2)}{(z-z')\sqrt{(x+W/2)^2 + (y+L/2)^2 + (z-z')^2}} \right] \end{array} \right] \quad (118)$$

The fields for a cuboid permanent magnet are the sum of the field from positive charge sheet at $z'=+h/2$ and a negative charge sheet at $z'=-h/2$.

Considering a point at the center of the face an infinitely long permanent magnet (so only one charge sheet) we should have $B_z = Br/2$. The B_z expression evaluates to 4π in this case, so we can conclude that $\mu_0 \sigma_m$ must equal $B_r/4\pi$.

$B_x(x, y, z) = \frac{B_r}{4\pi} \left[\begin{aligned} &+ \log \left[\frac{(y-L/2 + \sqrt{(y-L/2)^2 + (x+W/2)^2 + (z-H/2)^2})^{1/2} (y+L/2 + \sqrt{(y+L/2)^2 + (x-W/2)^2 + (z-H/2)^2})^{1/2}}{(y-L/2 + \sqrt{(y-L/2)^2 + (x-W/2)^2 + (z-H/2)^2})^{1/2} (y+L/2 + \sqrt{(y+L/2)^2 + (x+W/2)^2 + (z-H/2)^2})^{1/2}} \right] \\ &- \log \left[\frac{(y-L/2 + \sqrt{(y-L/2)^2 + (x+W/2)^2 + (z+H/2)^2})^{1/2} (y+L/2 + \sqrt{(y+L/2)^2 + (x-W/2)^2 + (z+H/2)^2})^{1/2}}{(y-L/2 + \sqrt{(y-L/2)^2 + (x-W/2)^2 + (z+H/2)^2})^{1/2} (y+L/2 + \sqrt{(y+L/2)^2 + (x+W/2)^2 + (z+H/2)^2})^{1/2}} \right] \end{aligned} \right]$	
$B_y(x, y, z) = \frac{B_r}{4\pi} \left[\begin{aligned} &+ \log \left[\frac{(x-W/2 + \sqrt{(x-W/2)^2 + (y+L/2)^2 + (z-H/2)^2})^{1/2} (x+W/2 + \sqrt{(x+W/2)^2 + (y-L/2)^2 + (z-H/2)^2})^{1/2}}{(x-W/2 + \sqrt{(x-W/2)^2 + (y-L/2)^2 + (z-H/2)^2})^{1/2} (x+W/2 + \sqrt{(x+W/2)^2 + (y+L/2)^2 + (z-H/2)^2})^{1/2}} \right] \\ &- \log \left[\frac{(x-W/2 + \sqrt{(x-W/2)^2 + (y+L/2)^2 + (z+H/2)^2})^{1/2} (x+W/2 + \sqrt{(x+W/2)^2 + (y-L/2)^2 + (z+H/2)^2})^{1/2}}{(x-W/2 + \sqrt{(x-W/2)^2 + (y-L/2)^2 + (z+H/2)^2})^{1/2} (x+W/2 + \sqrt{(x+W/2)^2 + (y+L/2)^2 + (z+H/2)^2})^{1/2}} \right] \end{aligned} \right]$	
$\bar{B}_z(x, y, z) = \frac{B_r}{4\pi} \left[\begin{aligned} &+ \tan^{-1} \left[\frac{(x-W/2)(y-L/2)}{(z-H/2)\sqrt{(x-W/2)^2 + (y-L/2)^2 + (z-H/2)^2}} \right] \\ &- \tan^{-1} \left[\frac{(x-W/2)(y+L/2)}{(z-H/2)\sqrt{(x-W/2)^2 + (y+L/2)^2 + (z-H/2)^2}} \right] \\ &- \tan^{-1} \left[\frac{(x+W/2)(y-L/2)}{(z-H/2)\sqrt{(x+W/2)^2 + (y-L/2)^2 + (z-H/2)^2}} \right] \\ &+ \tan^{-1} \left[\frac{(x+W/2)(y+L/2)}{(z-H/2)\sqrt{(x+W/2)^2 + (y+L/2)^2 + (z-H/2)^2}} \right] \\ &- \tan^{-1} \left[\frac{(x-W/2)(y-L/2)}{(z+H/2)\sqrt{(x-W/2)^2 + (y-L/2)^2 + (z+H/2)^2}} \right] \\ &+ \tan^{-1} \left[\frac{(x-W/2)(y+L/2)}{(z+H/2)\sqrt{(x-W/2)^2 + (y+L/2)^2 + (z+H/2)^2}} \right] \\ &+ \tan^{-1} \left[\frac{(x+W/2)(y-L/2)}{(z+H/2)\sqrt{(x+W/2)^2 + (y-L/2)^2 + (z+H/2)^2}} \right] \\ &- \tan^{-1} \left[\frac{(x+W/2)(y+L/2)}{(z+H/2)\sqrt{(x+W/2)^2 + (y+L/2)^2 + (z+H/2)^2}} \right] \end{aligned} \right]$	

Table 11: Magnetic flux density B around a cuboid permanent magnet of length L , width W , and height H , centered at the origin of Cartesian axes (x, y, z)

5.1.1 Notes on Computation

Sub-parts of these expressions have singularities at the surface of the magnet, but the full expression does come to a well-defined limit at points arbitrarily close to the surface of the magnet. For example, the terms in the tan expression go to infinity as z goes to $H/2$, but tan goes to pi as its argument goes to infinity, so (in this case) the full expression ends up being finite. To

compute these expressions as part of a simulation, it is efficient to check and see if the given point is on a boundary, and if so to perturb it slightly off the boundary.

The logs in the B_x and B_y expression can be combined for computational efficiency. It is attractive to attempt to combine the arctangents in the B_z expression, using the arctangent addition formula

$$\tan^{-1}(a) + \tan^{-1}(b) = \tan^{-1}\left(\frac{a+b}{1-ab}\right) + \begin{cases} \pi & a, b > 0 \\ \pi & a, b < 0 \\ 0 & \textit{otherwise} \end{cases} \quad (119)$$

However, code-profiling experiments using these expressions revealed that just giving gcc the full expression for B_z with 8 arctangents was more efficient than evaluating the arctangent addition expressions, either iteratively or in a tree, and also more efficient than evaluating the closed-form rational function version of the arctangent addition expression on pieces of four terms, then adding them. Because FPATAN instruction takes 270 machine cycles on the Intel IA-32 processor used to make the computation, divide about 17 cycles, and multiply about 4 cycles, this is puzzling --- unless gcc knows the arctangent addition formula and does term rewriting, which also seems unlikely.

It was profitable to precompute terms like $(x-W/2)$, the squares of these terms, and the expressions under radicals. All eight radical terms are used in computing the field in each of the three Cartesian directions, so precomputing them does indeed save execution time.

6 REFERENCES

- [1] Gershenfeld, N., *Fab: The Coming Revolution on Your Desktop --- From Personal Computers to Personal Fabrication*, Basic Books, New York, NY, 2005
- [2] e.g. Envision SR printer, 3D Systems, Rock Hill, SC
- [3] Fuller, S.B., Wilhelm, E.J., Jacobson, J.M., *Ink-jet printed nanoparticle microelectromechanical systems*, IEEE Journal of Microelectromechanical Systems, Volume 11, Issue 1, p. 54-60
- [4] B. A. Ridley, B. Nivi, and J. M. Jacobson, "All-inorganic field effect transistors fabricated by printing," *Science*, vol. 286, pp. 746-749, 1999.
- [5] Malone, E., Lipson, H., *Fab@Home: the personal desktop fabricator kit*, Rapid Prototyping Journal, Vol. 13, Issue 4, p. 245
- [6] Griffith, S., *Growing Machines*, MIT PhD Thesis, 2004
- [7] Rothmund, P., *Folding DNA to create nanoscale shapes and patterns*, *Nature*, 440, p. 297-302, 2006
- [8] Moore, G., *Cramming More Components onto Integrated Circuits*. *Electronics*, April 19, 1965.
- [9] Moore, G., *No Exponential is Forever...but We Can Delay 'Forever.'* Presentation at the IEEE International Solid-State Circuits Conference, February 10, 2003
- [10] Moore, G., *Visions of Technology*, Simon and Schuster, 1999
- [11] M. Bohr, et. al, *A 65 nm Logic Technology Featuring 35 nm Gate Lengths, Enhanced Channel Strain, 8 Cu Interconnect Layers, Low-k ILD and 0.57 μm^2 SRAM Cell*. Proceedings of the IEEE International Electron Devices Meeting, San Francisco, CA, 2004
- [12] Abelson, H., Allen, D., Coore, D., Hanson, C., Homsy, G., Knight, T.F., Nagpal, R., Rauch, E., Sussan, G.J., Weiss, R., *Amorphous Computing*, *Communications of the ACM*, Volume 43, Issue 5, pp. 74-82, May 2000
- [13] Butera, W. J. *Programming a paintable computer*. Massachusetts Institute of Technology Doctoral Thesis, 2002, UMI Order Number: AAI0804036.
- [14] Goldstein, S., Campell, J., Mowry, T., *Programmable Matter*, *Computer*, Vol. 38, No. 6, pp. 99-101, June 2005
- [15] M. Yim, W. Shen, B. Salemi, D. Rus, M. Moll, H. Lipson, E. Klavins, G.S. Ghirkjian, "Modular Self-Reconfigurable Robot Systems," *IEEE Robotics and Automation Magazine*, vol. 14, no. 1, Mar. 2007, pp. 43-52
- [16] Casal, A., Yim, M.H., Self-reconfiguration planning for a class of modular robots, *Proceedings of SPIE*, Issue 3839, p. 246-257, 1999
- [17] Chirikjian, G., Pamecha A., Ebert-Uphoff, I., Evaluating the efficiency of self-reconfiguration in a class of modular robots, *Journal of Robotic Systems*, Volume 13, No. 5, pp. 317-338, 1996
- [18] Pamecha, A., Ebert-Uphoff, I., Chirikjian G.S., *Useful Metrics for Modular Robot Motion Planning*, *IEEE Transactions on Robotics and Automation*, Vol.13, No.4, 1997.
- [19] Chen, I., Burdick, J., Determining Task Optimal Modular Robot Assembly Configurations, *Proceedings of the 1995 IEEE International Conference on Robotics and Automation*, Volume 1, p. 132-137, Nagoya, Japan, 1995
- [20] Werfel, J., Bar-Yam, Y., Rus, D., Nagpal, R., *Distributed Construction by Mobile Robots with Enhanced Building Blocks*, *Proceedings of the 2006 IEEE International Conference on Robotics and Automation, ICRA 2006*, May 15-19, 2006, pp. 2787-2794
- [21] Donald, B., Jennings, J., Rus, D., *Minimalism + Distribution = Supermodularity*, *Journal of Experimental and Theoretical Artificial Intelligence*
- [22] Kotay, K., Rus, D., Generic distributed assembly and repair algorithms for self-configuring robots, *Proceedings of the IEEE International Conference on Intelligent Robots and Systems, IROS 2004*, Vol. 3, p. 2362-2369, October 2004

-
- [23] Fukuda, T., Nakagawa, S., *Dynamically Reconfigurable Robotic System*, Proceedings of the IEEE International Conference on Robotics and Automation, April 1988, p. 1581-1586.
- [24] Yim, M., Zhang, Y., Roufas, K., Duff, D., Eldershaw, C., *Connecting and Disconnecting for Chain Self-Reconfiguration with Polybot*, IEEE/ASME Transactions on Mechatronics, Vol. 7, No. 4, December 2002
- [25] Castano, A., Behar, A., Will, P., *The Conro Modules for Reconfigurable Robots*, IEEE/ASME Transactions on Mechatronics, Vol. 7, No. 4., December 2002
- [26] Kurokawa, H. Kamimura, A. Yoshida, E. Tomita, K. Murata, S. Kokaji, S. *Self-reconfigurable modular robot (M-TRAN) and its motion design*, International Conference on Control, Automation, Robotics, and Vision, 2002, Dec 2-5, 2002
- [27] Kotay, K., Rus, D., *Locomotion versatility through self-reconfiguration*, Robotics and Autonomous Systems, Vol. 26, Issues 2-3, Feb 28 2006, p. 217-232
- [28] Rus, D., Vona, M., *A basis for self-configuring robots using crystal modules*, Proceedings of the International Conference on Intelligent Robots and Systems, IROS 2000, Vol. 3, p. 2194-2202
- [29] Unsal, C., Kilccote, H., Khosla, P.K., *A Modular Self-Reconfigurable Bipartite Robotic System: Implementation and Motion Planning*, Autonomous Robots, Vol 10, No. 1, January 2001
- [30] Zykov, V., Mytilinaios, E., Adams, B., Lipson, H., *Self-Reproducing Machines*, Nature, Volume 435, May 12, 2005, pp. 163-164
- [31] Shen, W., Krivokon, M., Chiu, H., Everist, J., Rubenstein, M., Venkatesh, J., *Multimode Locomotion via Superbot Robots*, Proceedings of the 2006 IEEE International Conference on Robotics and Automation, Orlando, Florida, May 2006
- [32] White, P.J., Kopanski, K., Lipson, H., *Stochastic self-reconfigurable cellular robotics*, Proceedings of the IEEE International Conference on Robotics and Automation, ICRA 2004, Volume 3, p. 2888-2893, May 2004
- [33] Griffith, S., Goldwater, D., Jacobson, J., *Self-replication from random parts*, Nature, Vol 437, p. 636, September 2005
- [34] Zykov, V., Lipson, H., *Experiment Design for Stochastic Three-Dimensional Reconfiguration of Modular Robots*, Proceedings of IROS 2007
- [35] B. Kirby, B. Aksak, J. Hoburg, T. Mowry, P. Pillai, "A Modular Robotic System Using Magnetic Force Effectors," IROS 97
- [36] Gilpin, K., Kotay, K., Rus, D., *Miche: Modular Shape Formation by Self-Dissassembly*, IEEE International Conference on Robotics and Automation, April 10-14, 2007, Roma, Italy
- [37] Warneke, B.A. Scott, M.D. Leibowitz, B.S. Lixia Zhou Bellew, C.L. Chediak, J.A. Kahn, J.M. Boser, B.E. Pister, K.S.J. *An autonomous 16 mm/sup 3/ solar-powered node for distributed wireless sensor networks*. Proceedings of IEEE Sensors, 2002, Volume 2, Page 1510-1515, ISBN 0-7803-7454-1
- [38] Drzaic, P., Chiang, A., Stewart, R., Hermanns, A., Shi, Y., Jacobsen, J. *Plastic-film displays with NanoBlock IC drivers integrated by a fluidic self-assembly process*. Journal of the Society for Information Display, 2003, Volume 11, Issue 1, pp. 81-87
- [39] Senturia, S.D., *Microsystem Design*. Kluwer Academic Publishers, 2001.
- [40] Gray, P.R., Meyer, R.G., *Analysis and Design of Analog Integrated Circuits*, John Wiley and Sons, 1993
- [41] Mills, A.F. *Basic Heat and Mass Transfer*. 2nd Edition, Prentice Hall, 1998
- [42] Tarascon, J.M., Armand, M. *Issues and challenges facing rechargeable lithium batteries*. Nature, p.359, Volume 414, November 2001,
- [43] Duracell. *Zinc-Air Technical Bulletin*. Boston: The Gillette Company, 2004.
- [44] Bates, J.B., Gruzalski, G.R., Luck, C.F., *Rechargeable solid state lithium microbatteries*. Proceedings of the 6th IEEE workshop on MEMS, Fort Lauderdale, FL, 1993, pp. 82-86
- [45] Epstein, A.H, Senturia, S.D., Anathasuresh, G., Ayon, K. *Power MEMS and microengines*, Technical Digest of the International Conference on Solid-State Sensors and Actuators, 1997

-
- [46] Madou, M. *Fundamentals of Microfabrication*. 2nd edition. CRC Press, 2002
- [47] Blanchard, J., Henderson, D., Lal, A., *Nuclear Microbattery for MEMS Devices*. Project report submitted to U.S. Department of Energy
- [48] MSDS, ⁶³Ni
- [49] MSDS, ³H
- [50] Hall, E., *Radiobiology for the Radiologist*. 4th Edition, Lippincott Williams & Wilkins, 1993
- [51] Shapiro, J., *Radiation Protection: a guide for scientists and physicians*. Cambridge: Harvard University Press, 1990
- [52] Guo, H., Lal, A., *Nanopower Betavoltaic Microbatteries*. TRANSDUCERS: 12th International Conference of Solid-State Sensors, Actuators, and Microsystems, Volume 1, p. 36-39, June 2003.
- [53] Wang, S. Li, S., Du, Y., Xu, B., Li, L., Zhu, Y. Size Limit of Zinc Nanoparticles. The Preliminary Program for the Spring National Meeting on Engineered Particles or Engineered Nanoparticle Structures, Orlando F.L., 2006
- [54] Paradiso, J. A. and Starner, T. 2005. *Energy Scavenging for Mobile and Wireless Electronics*. IEEE Pervasive Computing 4, 1 (Jan. 2005), 18-27.
- [55] Virtuani, A.; Lotter, E.; Powalla, M.; Rau, U.; Werner, J.H. *Highly resistive Cu(In,Ga)Se₂ absorbers for improved low-irradiance performance of thin-film solar cells*. Thin Solid Films, 2004, Vol. 451-452, pp. 160-165
- [56] M. Green et. al, *Solar Cell Efficiency Tables (Version 18)*, Progress in Photovoltaics: Research and Applications, Volume 9, Wiley/Interscience, 2001
- [57] Ogale, S.B. et. al, *High Temperature Ferromagnetism with a Giant Magnetic Moment in Transparent Co-doped SnO₂δ*. Physical Review Letters, Vol. 91, No.7, August 2003
- [58] Park, J., Allen, M. *Ultralow-Profile Micromachined Power Inductors With Highly Laminated Ni/Fe Cores: Applications to Low-Megahertz DC-DC Converters*. IEEE Transactions on Magnetics, Vol. 39, No. 5, September 2003
- [59] Liang-Liang, X., Kumar, P.R., *A network information theory for wireless communications: scaling laws and optimal operation*. IEEE Transactions on Information Theory, Volume 50, Issue 5, pp. 748-767
- [60] Lee, T.H, *CMOS RF: no longer an oxymoron*. Proceedings of the 19th Gallium Arsenide Integrated Circuit Symposium, pp. 244-247, 1997.
- [61] Camilleri, N., Costa, J., Lovelace, D., Ngo, D., *Silicon MOSFETs, the microwave device technology for the 90s*. 1993 IEEE International Microwave Symposium Digest p. 545, 1993.
- [62] Lee, T.H., *The Design of CMOS Radio-Frequency Integrated Circuits*. Cambridge: Cambridge University Press, 1998
- [63] Leibowitz, B., Boser, B., Pister, K.S.J., *CMOS smart pixel for free-space optical communication*. Proceedings of SPIE – Volume 4306: Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications II, May 2001, pp. 308-318
- [64] Sunderarajan, M., Hershenson, M., Boyd, S., Lee, T.H., *Simple Accurate Expressions for Planar Spiral Inductances*. IEEE Journal of Solid-State Circuits, Vol. 34, No. 10, October 1999
- [65] Lee, Y., *Antenna Circuit Design for RFID Applications*. Microchip Technology Corp, Chandler, AZ, AN710, 2003
- [66] Balanis, C.A., *Antenna Theory: Analysis and Design*. 2nd Edition, John Wiley and Sons, 1996
- [67] Lin, J., Sugavanam, A., Guo, X., Li, R., Brewer, J., Kenneth, K.O., *Integrated Antennas on Silicon Substrates for Communication Over Free Space*. IEEE Electron Device Letters, Vol. 25, No. 4, April 2004
- [68] Porret, A.S., Melly, T., Python, D., Enz, C.C., Vittoz, E.A. *An ultralow-power UHF transceiver integrated in a standard digital CMOS process: architecture and receiver*. IEEE Journal of Solid State Circuits, Volume 36, Issue 3, pp. 452-466, March 2001

-
- [69] Doran, C., Emami, S., Sobel, D., Niknejad, A., Broderson, R., *Design Considerations for 60 GHz CMOS Radios*. IEEE Communications Magazine, pp. 132-140, December 2004.
- [70] Advanced RISC Machines Corp., *ARM7TDMI: ARM 32-bit RISC core with 16-bit system costs*. <http://www.arm.com/products/CPUs/ARM7TDMI.html>
- [71] M. Bohr, et. al, *A 130nm Generation Logic Technology Featuring 70 nm Transistors, Dual Vt Transistors, and 6 layers of Cu Interconnects*. Proceedings of the IEEE International Electron Devices Meeting, San Francisco, CA, 2000
- [72] Yanagiya, N., Matsuda, S., Inaba, S., Takayanagi, M. Mizushima, I. Ohuchi, K. Okano, K. Takahasi, K. Morifuji, E. Kanda, M. Matsubara, Y. Habu, M. Nishigoori, M. Honda, K. Tsuno, H. Yasumoto, K. Yamamoto, T. Hiyama, K. Kokubun, K. Suzuki, T. Yoshikawa, J. Sakurai, T. Ishizuka, T. Shoda, Y. Moriuchi, M. Kishida, M. Matsumori, H. Harakawa, H. Oyamatsu, H. Nagashima, N. Yamada, S. Noguchi, T. Okamoto, H. Kakumu, *65 nm CMOS technology (CMOS5) with high density embedded memories for broadband microprocessor applications*, Technical Digest of the IEEE Electron Devices Meeting, IEDM 2002, p. 57, 2002.
- [73] Keeney, S., *A 130nm Generation High Density Etox™ Flash Memory Technology*. Technical Digest of the IEEE International Electron Devices Meeting, Washington D.C., pp. 2.5.1-2.5.4, Washington, DC, 2001
- [74] Atwood, G., *Future Directions and Challenges for ETox Flash Memory Scaling*. IEEE Transactions of Device and Materials Reliability, Vol. 4, No.3, September 2004
- [75] Zahn, M., *Electromagnetic Field Theory: A Problem-Solving Approach*, Krieger Publishing Company, Malabar, FL, Reprint Edition with corrections, 2003
- [76] Benhammou, D., *Large Scale Electrostatic Actuation*, Masters Thesis, MIT
- [77] Gutmann, R.J., Lu, J.Q., Kwon, Y., McDonald, J.F., Cale, T.S. *Three-dimensional (3D) ICs: a technology platform for integrated systems and opportunities for new polymeric adhesives*. First International IEEE Conference on Polymers and Adhesives in Microelectronics and Photonics, pp. 173-180, Potsdam, Germany, 2001
- [78] Yang et. al, *Monolithic integration of III-V optical interconnects on Si using SiGe virtual substrates*. Journal of materials science: Materials in electronics, Vol. 13, Issue 7, Page 377, 2002
- [79] Pollentier, I. et. al, *Epitaxial lift-off GaAs LEDs to Si for fabrication of opto-electronic integrated circuits*, IEEE Electronics Letters, Vol 26, No. 3, pp. 193-4, February 1990
- [80] Personal Communication with David Tennenhouse, who is the Director of Research for Intel Corporation, September 2004.
- [81] Jacobson, J., Gershenfeld, N., Butera, B., *The V1.0 'Pushpin' Particle 2-D Self-Assembling Display Array*. DAPRA / Natick Technical Report, April, 2004
- [82] Lifton, J., Seetharam, D., Broxton, M., and Paradiso, J. A. Pushpin Computing System Overview: A Platform for Distributed, Embedded, Ubiquitous Sensor Networks. In *Proceedings of the First international Conference on Pervasive Computing* (August 26 - 28, 2002). F. Mattern and M. Naghshineh, Eds. Lecture Notes In Computer Science, vol. 2414. Springer-Verlag, London, 139-151.
- [83] Basak, A., *Permanent-magnet DC linear motors*, Clarendon Press, Oxford, 1996
- [84] Slocum, A., *Precision Machine Design*, Society of Manufacturing Engineers, Dearborn, Michigan, 1992
- [85] *ADXL320 Data Sheet*, Analog Devices, Norwood, MA
- [86] Niino, T. Higuchi, T. Egawa, S., Dual excitation multiphase electrostatic drive, Industry Applications Conference, 1995. Thirtieth IAS Annual Meeting, IAS '95., Conference Record of the 1995 IEEE, 8-12 Oct 1995, Volume 2, p. 1318-1325, Orlando, FL
- [87] Horowitz, P., Hill, W., *The Art of Electronics*, Cambridge University Press, 1989
- [88] Reiche, M., Wagner, G., *Wafer Thinning: Techniques for Ultra-thin Wafers*, Advanced Packaging, March 2003
- [89] Personal Communication, William Butera, Intel Corporation

-
- [90] *Method of thinning integrated circuits received in die form*, U.S. Patent 6013534, January 11, 2000
- [91] Lee, R.A., Whittaker, D.R., *Laser created silicon vias for stacking dies in MCMs*, Electronics Manufacturing Technology Symposium, 1991., Eleventh IEEE/CHMT International
- [92] Burkett, S.L, Qiao, X., Temple, D., Stoner, B., McGuire, G., *Advanced processing techniques for through-wafer interconnects*, J. Vac. Sci. Technol. B 22(1), Jan/Feb 2004
- [93] Davey, K., Klimpke, B., *Computing Forces on Conductors in the Presence of Dielectric Materials*, IEEE Transactions on Education, Vol. 45, No. 1, February 2002
- [94] Mark Yim John Lamping Eric Mao J. Geoffrey Chase, *IC-processed electrostatic micro-motors*, Electron Devices Meeting, 1988. IEDM '88. Technical Digest., International, p. 666-669, 1988, San Francisco, CA
- [95] Emmanouel Hourdakias, a □ Brian J. Simonds, and Neil M. Zimmerman, *Submicron gap capacitor for measurement of breakdown voltage in air*, Rev. Sci Instrum. 77 034702 (2006)
- [96] Egawa, S., Niino, T., Higuchi, T., *Film actuators: Planar, electrostatic surface-drive actuators, Micro Electro Mechanical Systems*, 1991, MEMS '91, Proceedings. 'An Investigation of Micro Structures, Sensors, Actuators, Machines, and robots' January 1991, Nara, Japan
- [97] M. Bohr, et. al, *A 65 nm Logic Technology Featuring 35 nm Gate Lengths, Enhanced Channel Strain, 8 Cu Interconnect Layers, Low-k ILD and 0.57 μm^2 SRAM Cell*. Proceedings of the IEEE International Electron Devices Meeting, San Francisco, CA, 2004
- [98] *ARM7TDMI core data sheet*, ARM, Inc., <http://www.arm.com/products/CPUs/ARM7TDMI.html>
- [99] Butler, Z., Kotay, K., Rus, D., Tomita, K., *Generic decentralized control for a class of self-reconfigurable robots*, Robotics and Automation, 2002. Proceedings. ICRA '02. IEEE International Conference on, Volume 1, pp. 809-816