

Volume Mount Devices

by

Alan Han

B.S. Electrical Engineering, Purdue University (2018)

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

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2025

© 2025 Alan Han. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Authored by: Alan Han
Program in Media Arts and Sciences
May 16, 2025

Certified by: Prof. Neil Gershenfeld
Director, MIT Center for Bits and Atoms
Thesis Supervisor

Accepted by: Tod Machover
Academic Head, Program in Media Arts and Sciences

Volume Mount Devices

by

Alan Han

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning,
on May 16th, 2025 in partial fulfillment of the requirements for the degree of

Master of Science

Abstract

As Moore’s Law ends and AI demands increasingly tax our climate and resources, the limitations of two-dimensional electronics integration have become critical bottlenecks. Surface-mount devices (SMDs) remain entrenched in industry practice despite being insufficient for today’s computing challenges and sustainability needs. This thesis introduces the volume mount device (VMD), a three-dimensional electronics packaging standard that bypasses the traditional die-to-server stack while offering a scalable, reversible framework inspired by natural ecosystems’ circularity.

The VMD approach embeds both electrical function and mechanical structure into modular elements that assemble freely in 3D space. Rather than building circuits on planar PCBs, this system constructs functional circuits by linking components into a self-constraining lattice architecture. My current implementation leverages existing supply chains by incorporating SMD components on small tile PCBs, while establishing a pathway toward eventually replacing SMDs at the IC packaging level.

I developed a hybrid assembly system combining 3D printing and pick-and-place automation to build multi-layered electronic assemblies efficiently. Where prior work achieved only tens of parts at hundreds of components per hour (CPH), my system demonstrates automated assembly of hundreds of integrated elements at approximately 1000 CPH. I evaluate various geometric configurations, assess performance overhead compared to conventional approaches, and develop cost-effective, self-aligning connector interfaces for reliable joints—creating a foundation for electronics systems that can be assembled, disassembled, and reassembled as needed while improving resilience against supply chain disruptions.

Thesis supervisor: Prof. Neil Gershenfeld

Title: Director, MIT Center for Bits and Atoms

Volume Mount Devices

by

Alan Han

This thesis has been reviewed and approved by the following committee members:

Prof. Neil Gershenfeld

Director, Center for Bits and Atoms
Professor, Media Arts and Sciences
Massachusetts Institute of Technology

Prof. Ahmad Bahai

Senior Vice President and Chief Technology Officer, Texas Instruments
Professor of the Practice, EE
Massachusetts Institute of Technology

Prof. Joseph Jacobson

Associate Professor of Media Arts and Sciences
Massachusetts Institute of Technology

Acknowledgements

Thank you to everyone who has supported me along the way, I am beyond grateful and could not have made this journey without you.

I'd like to start by thanking Neil, for creating this amazing space with amazing people and tools, and challenging me to rethink what I'm capable of. Thank you to my readers Joe and Ahmad, for your support and insight as I developed this thesis.

Thank you to the CBA of past, present, and future, as well as visitors, UROPs, and friends of the lab, for making the basement a lively community: Quentin, Nik, Jake, Alfonso, Miana, Kat, Jiaming, Marcello, Shravika, Kerdeem, Yohan, Ric, Leo, Cedric, Dave, Zach, Will, Amira, Camron, Alex, Erik, Sara, Paul, Scott, Nicole, Eyal, Jack, Francisco, and Vlasta.

Thank you Miana, for convincing me to apply, for which I'm glad I did. Thank you to the usual suspects, Quentin and Nik, and Jake, and Alfonso, who always seem to look out for me. Thank you Erik, for providing the architectural foundation for this project and giving me an excuse to decorate my keyboard with new symbols, and for being patient with my many questions. Thank you Zach, Will, and Jonathan, for pioneering discrete electronics in the CBA and for providing invaluable reference points. I stand on your shoulders.

To my Media Lab community: thank you Retroentomologists and Followers; especially Michael, Ana, Char, Nathan, Lucy, Eleanor, Sarah, Gauri, Isabella, Perry, and Lancelot. Our cohort kept me cohesive and you put the magic in the media lab for me. Sarah, Otto, and Dave, thank you for our game-of-thrones caliber commander games. Cedric, thank you for always being down to hack something and organizing the Shenzhen trip for which I am grateful for.

Thank you Foppers and Triangles back home, especially Amy, Cecilia, Emily, Satyo, and Ed, for keeping me sane, and keeping me grounded. Thank you James, for your lifelong friendship. To my fellow Boilermakers, thank you Ryan and Nathan for the diving and supreme commanding. Thank you Charles and David, for formally solving problems with me and blowing up my phone with 500 notifications about naval doctrine and kiminonawa.

Thank you to everyone at Molex, and especially Vic for getting me obsessed about finding new ways to package electronics in the first place.

Thank you CBA sponsors and grant providers, for supporting our work and sharing our excitement for where it goes. Thank you CBA staff for keeping things running, John, Joe, Hannah, Kara, Candace, Dan, Marissa, and Sherry.

And last but not least, thank you to my family, Mom, Dad, and Angel, for your unconditional love, support, and snacks.

Table of Contents

1 Introduction	17
1.1 Motivation	17
1.2 Electronics Packaging	18
1.2.1 Printed Circuit Boards	19
1.2.2 Integrated Circuit Packaging	19
1.2.3 Heterogeneous Integration	19
1.3 Supply Chain Challenges and Climate Implications	20
1.4 Accessible Modularity	20
1.5 Related Work	21
1.5.1 Project Tinkertoy	21
1.5.2 Electronic Digital Materials	21
1.6 Thesis Contributions and Overview	25
2 Volume Mount Devices (VMDs)	28
2.1 Geometry Overview	28
2.2 4Hx Family	31
2.2.1 4H	31
2.2.2 4Hp	31
2.3 4Bx Family	37
2.3.1 4B	37
2.3.2 4BI	39
3 Circuit Assembly System	43
3.1 Assembly	43
3.1.1 Setup	43
3.1.2 Operation	45
3.2 Motion System	45
3.3 Feeders	45
3.3.1 Passive Tray Feeders	45
3.3.2 Adding Depth	46
3.3.3 Alternative Feeders	48
3.3.4 Resetting Jobs	48
3.4 Substrates	48
3.4.1 4Hp Macro-tile	48
3.4.2 4BI Template	49
3.4.3 Material Choice	51
3.5 End Effector	51
3.5.1 Claw Gripper	51
3.5.2 Vacuum Nozzle	52
3.5.3 Sealing the Vacuum	54
3.5.4 Nozzle Geometry	54
3.5.5 Tile Surface Planarity and Area	54
3.6 Error Correction	54
3.6.1 Overconstraint and Compliance	57
3.7 Comparative Assembly Rates	57

3.7.1 Possible Parallelization Pathways	59
4 Testing and Performance Overhead	62
4.1 Joint Reliability	62
4.1.1 4Hp H-eon Elements	62
4.1.2 4BI Pad-to-Pad	62
4.1.3 4BI Pad Geometry	62
4.2 Circuit Performance	70
4.2.1 Ring Oscillators	70
4.2.2 Logic Gates	70
4.2.3 Full Adders	70
4.3 Scaling Performance Projections	73
4.3.1 Capacitance	73
4.3.2 Inductance	75
4.3.3 Resonance Frequency	76
4.3.4 Capacitive Switching Energy	76
4.3.5 Projections from 10mm down to 10nm	76
5 4BI Connector (4BIc)	82
5.1 Electrical Connectors in a Day (eCiD)	82
5.1.1 Assembly Techniques	82
5.1.2 Housings	84
5.1.3 Leadframes	84
5.1.4 Process Registration	87
5.2 4BIc Geometry	88
5.2.1 4BIc Process Steps Overview	89
5.2.2 Panelized PCBs	89
5.2.3 Leadframe Fabrication	90
5.2.4 Ejecting Trash	91
5.2.5 Form 1 Tooling for Leadframe	91
5.2.6 Formed Leadframe	93
5.2.7 Form 2 Tooling for Leadframe and Pcb	93
5.2.8 Tool Singulation of Pcb Panel	93
5.2.9 Advancing VMDs	96
6 Conclusion	99
6.1 Future Work	99
6.1.1 Finishing 4BIc	99
6.1.2 Assembly Improvements	100
6.1.3 Geometric Exploration for Enhanced Reliability	100
6.1.4 4Cx Family - 4C, 4Ci, 4Ce, 4Cs	101
6.1.5 uVMDs and nVMDs	101
6.1.6 Mini-Assembly Farms for Smaller Elements	104
6.1.7 Super-DICE 2.0	104
6.2 Impact	105
Bibliography	107

List of Figures

1.1 Hierarchy of electronics packaging levels. Level 0 represents the bare die (transistors). Level 1 is the integrated circuit (IC) package. Level 2 consists of the substrate (e.g., PCB). Level 3 refers to the full system.	18
1.2 Different Heterogeneous Integration approaches	19
1.3 Digital electronic materials history, Neil Gershenfeld	21
1.4 Basic Bulk Materials	22
1.5 Component Materials	23
1.6 Integrated Materials	24
1.7 Component level, in-between basic and integrated	24
2.1 4BIc0 circuit assemblies: 4BIc0 8-layer mechanical benchmark (254/256 tiles) automatically assembled in ~14min (a); 4BIc0 2-layer stackable full-adder circuit (146 tiles) manually assembled in ~55min (b)	28
2.2 A variety of geometry iterations	29
2.3 Fabricated geometry families	29
2.4 4H dimensions: the tile (a), the H connector (b), and the lattice (c)	32
2.5 4H overview	33
2.6 4Hp: the tile (a), the connector (b), and the lattice (c)	34
2.7 Top: H-eon, from cad to fab, Bottom: 4Hp tile, from cad to fab	35
2.9 left, basic 6-tile library; right, implemented power, ground, and signal mapping for tiles	36
2.10 4Hp Batch 009 tiles	36
2.11 Hierarchical assembly of tiles and macro-tiles, which enable error-correction at larger, block-sized scales	36
2.12 4B design and assembly tests	37
2.13 4BI dimensions: 4BI tile (a), 4BI lattice pitch (b), 4BI template top down (c), template corner features (d), and side profile of flexures (e)	39
2.14 initial 4BI geometry concept and prototypes (Neil Gershenfeld)	40
2.15 4BI tiles	41
2.16 4BI tiles in gridfinity bins for storage (a), and tiles loaded in a feeder for assembly (b)	41
3.1 Benchmark tests	44
3.2 The assembler picking and placing from 4BIc0 feeders	44
3.3 Basic PnP operation (without computer vision): picking (a), moving (b), placing (c)	45
3.4 LumenPnP	46
3.5 Parametric feeder design in CadQuery: 30 element capacity 3x10_feeder (a), 60 element capacity 3x20_feeder (b), 60 element capacity 6x10_feeder (c); evolution of feeder concepts for 4Hp, from trays to magazine dispensers (d)	47
3.6 4Hp macro tiles as substrates	49
3.7 Parametric template design in Cadquery: 9x9x10 template, 320 element capacity (a), 9x9x3 template, 96 element capacity (b), 17x17x10 template, 1280 element capacity (c)	49
3.8 Templates in use for various tasks: early rigid template with poor tolerance (a),	

high-temp reflow variation from soldering tests (b), matured design using Form 4 Black V5 (c)	50
3.9 Evolution of tile gripper designs for 4Hp	52
3.10 Passive gripper end effector for 4Hp H-eon connector	53
3.11 4BI error states: failure to rotate, tile never enters template (a), tile not completely flush; opportunity to correct but may interfere with adjacent placements (b), tile placed on its side; tile enters template but is invalid and affects adjacent tiles (c)	55
3.12 Vacuum end effector: Tile tilting due to uneven solder, center support (a), CAD of end effector supporting tile on edges, even force distribution (b)	56
3.13 SlimStack connectors from Tiny-DICE, broken from overconstraint [42]	57
3.14 Will Langford's scaling throughput graph [47]	57
4.1 4BI experiments to try and solve intermittency: Spring-loaded compression cap (a), solder and reflow (b,c), experimenting with plating finishes (d), introducing contact geometry via leadframes (e), mechanical decoupling of cable strain	63
4.2 4BI pitch	64
4.3 At 5.655mm pitch, 1mm x 1mm pads (4BI-b001, b002) could easily result in only ~25% pad overlap with underconstrained rotational error. 1.5mm x 1.5mm pads (4BI-b003+) increases this to ~50%, with little downside. Reducing the pitch would improve this even more	64
4.4 4BIc contact improvements	65
4.5 Test jigs for single tile pad-to-pad resistance tests: CAD model of test jig (a), soldering stranded wire to pyralux, solder joint to tile (b), test jigs fabricated for each of the 4 test finishes (c), test jigs in their closed state, ready for testing (d).	66
4.6 pad-to-pad resistance vs compression force test setup	67
4.7 4-wire test setup, measuring $R_{\text{PadJunction}}$	67
4.8 Plating finish test results, resistance vs compression load	68
4.9 Continuity test setup: circuit assembly (a), setup (b), circuit in dice-cad (Erik Strand) (c)	69
4.10 2MHz resonant frequency 4Hp ring oscillator circuit: ring oscillator schematic (a), 4Hp ring oscillator circuit (b), oscilloscope shot of circuit ringing (c), 4BIc0 ring oscillator design in dice-cad (d)	70
4.11 4Hp logic implementations: 4Hp diode AND gate (a), diode AND gate schematic (b), 4Hp nmos NAND gate (c), nmos NAND gate schematic (d)	71
4.12 Adder circuits: Full-adder schematic (a), 4Hp Full-adder implementation (Shravika Pendyala) (b), 4BIc0 half-adder (c), 4BIc0 full-adder (d)	71
4.13 ~146 tile 4BIc0 full-adder circuit assembled (a), and in test-setup (b)	72
4.14 modeling parasitic capacitance between two vertically neighboring tiles (in green), which may have different, partially overlapping routing patterns	74
4.15 modeling parasitic inductance loop through the smallest possible routing loop between 2x2 tiles	75
4.16 Parasitic Parameters vs Tile Size ($d_1 = d_2 = 0.33 * s$, Two-Dielectric Model): Capacitance vs Tile Size (a), Inductance vs Tile Size (b), LC Product vs Tile Size (c), Capacitive Energy vs Tile Size (d)	77
4.17 Parasitic Parameters vs Tile Size, Sweep over Coverage ($d_1 = d_2 = 0.33 * s$, Two-Dielectric Model): Capacitance vs Tile Size (a), Inductance vs Tile Size (b), LC	

Product vs Tile Size (c), Capacitive Energy vs Tile Size (d)	78
4.18 Parasitic Parameters vs Tile Size, Sweep over $d1 = d2 = [0.1-0.33] * s$, (Two-Dielectric Model): Capacitance vs Tile Size (a), Inductance vs Tile Size (b), LC Product vs Tile Size (c), Capacitive Energy vs Tile Size (d)	79
4.19 Routing considerations: parallel layer 1 (a), parallel layer 2 (b), crossed layer 1 (c), crossed layer 2 (d)	80
5.1 Connectors fabricated using a combination of desktop fabrication tools	83
5.2 Assembly techniques	83
5.3 Printing supports for small connector housings	85
5.4 Backside channels to alleviate cupping	86
5.5 Leadframe flow for the IDC	86
5.6 Form tool for the IDC	87
5.7 Fabricated tools for the IDC	87
5.8 4BIc dimensions: bare 4BIc0 tile dimensions (a), 4BIc tile with contacts (b), side view of the contact geometry (c)	88
5.9 4BIc process steps: (a) PCB 5x5 v-cut panel of sot323 tiles, (b) Fiber laser-cut leadframe, (c) Eject trash for solder rivets later, (d) Form 1: 90 degree bend, (e) Eject leadframe, (f) Formed leadframe 1, (g) Flip over, (h) Register against PCB for reflow step, (i) Form 2: 60 degree bend, create compliant contacts, (j) Formed leadframe 2, (k) Laser Singulation: remove leadframe tie-bars, (l) Tool Singulation: shear PCB v-cuts	89
5.10 4BIc assembled into a circuit	89
5.11 Laser-marking soldermask for differentiating sides	90
5.12 0.127mm C510 leadframe fabrication with the F1 Ultra	91
5.13 Tool assembly process: the print is post-cured in the Form Cure to maximum strength properties (a), parts, particularly the mating clamp, are checked for clear channels (b), tools are installed into their receptacles (c), the top half is then loaded onto the linear guides (d), the tool is loaded into the arbor press (e), and then the press is actuated and the leadframe has been modified, whether ejected/formed/singulated	92
5.14 The Form tooling process	92
5.16 Formed leadframe, with revised geometry solving earlier problems	94
5.17 The second form operation, which uses slanted tooling features to form the final leaf spring contacts	95
5.18 Tile singulation tool	95
5.19 Singulation troubleshooting: The first singulation left a pattern that correlates with the first wave of punches (a), the rest were stuck in the clamp (b). Increasing the opening of each of the windows resolved the issue. After ~20 singulations, edge quality seems to be fairly consistent (e).	96
6.1 4Cx family of geometries	102
6.2 ~55nm gold nano cubes undergoing contact printing for precise, scalable, and pristine particle patterning, [64]	102
6.3 4Cx VMD prototypes (with pockets for integrating components) printed on the UpNano One 2-photon printer	103
6.4 Super-DICE 2.0: AQFP buffer device designed by Camron Blackburn et al [69] (a), Concept layout for Super-DICE 2.0, Camron Blackburn's sketch (b)	105

List of Tables

2.1 Geometry feature comparison	30
3.1 Pick and Place Machine CPH Comparison spanning multiple categories	58
3.2 Selected PnP machines for projections	59
3.3 Parallelization pnp projections	59
4.1 Parasitic resistance measurements for each sample	68
4.2 the resulting Normal Force Window for all 4 samples	69

1 Introduction

1.1 Motivation

Moore’s law originated with Gordon Moore, who in 1965, observed that the number of transistors on an integrated circuit were doubling every two years. Later revised to every 1.5 years, Moore projected that transistors were on track to continue this trend for the next 10 years.

The practical implementation of Moore’s law was initially enabled by Dennard scaling, a principle formulated by Robert Dennard in 1974. Dennard scaling suggested that as transistors shrink, power density remains constant, allowing smaller transistors to operate at lower voltage, higher speeds, and with better power efficiency. This complementary principle meant that not only could more transistors fit on a chip (following Moore’s law), but they could also operate more efficiently, creating a self-sustaining cycle of performance improvement.

In fact, for the past 50+ years, Moore’s law has largely dominated the exponential performance increase seen by the computing industry, but due to the breakdown of Dennard scaling around 2006 [1] and physical limits, many have proclaimed Moore’s law dead [2]; [3]; [4]. Without Dennard scaling’s benefits, increased transistor density no longer automatically translates to proportional performance and efficiency gains. Transistors are still being improved [5], but no longer at the pace of Moore’s law, and with a higher price tag [6].

The end of Moore’s law has brought about demand for alternative solutions to continue increasing performance in lithography. This has placed a lot of focus on new approaches such as the Gate All Around FET (GAAFET) [7], back-side power delivery to alleviate on-device wiring density [8], and High-NA Extreme Ultra-Violet (EUV) lithography [9]; [10]; [11].

In particular, chiplets and heterogeneous integration (HI) [12]; [13] are a packaging innovation that enables designers to bring critical circuits closer together and fab designs on process nodes suitable for each chiplet, rather than compromising performance by using one node for the entire die, which can be challenging and expensive; power, compute, and MEMs all have different requirements, and different nodes are better for optimized for different applications.

While HI is now considered a mainstream advanced packaging technique that has allowed the industry to continue scaling performance, many of the viable implementations are still limited in the number of chiplets they can stack and interconnect; true 3d packaging without restriction is still out of reach, economically and technically. And while the semiconductor industry has access to HI, the rest of the electronics packaging stack has not been similarly improved, which reduces the effectiveness and accessibility

of gains made in the semiconductor industry.

In this thesis, I focus on ground work developing the Volume Mount Device (VMD), a term initially coined by Ahmad Bahai and Neil Gershenfeld. This new packaging standard aims to eventually replace SMD and enable true 3d heterogeneous integration, refactoring the entire electronics stack instead of a small portion of it.

1.2 Electronics Packaging

Electronic packaging refers to the design and production of enclosures for electronic devices ranging from individual semiconductor devices up to complete systems. Packaging must address critical constraints including protection from mechanical damage, cooling, electromagnetic interference (EMI), and electrostatic discharge.

The conventional electronics packaging (EP) paradigm describes how electronics are physically packaged across multiple levels, from semiconductor devices to complete systems:

Packaging Hierarchy Large Electronic System

Level	Description of Interconnection
0	Intraconnections on the chip
1	Chip-to-package interconnections to form IC package
2	IC package to circuit board interconnections
3	Circuit board to rack; card-on-board packaging
4	Wiring and cabling connections in cabinet




Figure 1.1: Hierarchy of electronics packaging levels. Level 0 represents the bare die (transistors). Level 1 is the integrated circuit (IC) package. Level 2 consists of the substrate (e.g., PCB). Level 3 refers to the full system.

This EP structure describes the physical hierarchy of electronic systems, while the OSI model describes the digital communication hierarchy. Interconnects between these levels form their own hierarchy, often denoted as “half-levels” (0.5, 1.5, 2.5), representing the critical interfaces between major electronics packaging levels.

While this established EP hierarchy has enabled tremendous growth in computing capability, we’re now facing challenges with greater compute demands, power constraints, thermal limitations, and efficiency requirements. The semiconductor industry has responded with advanced packaging approaches, but these innovations have primarily focused on level 0 and 1, with limited transformation at higher levels.

1.2.1 Printed Circuit Boards

Printed circuit boards have evolved dramatically since Paul Eisler created the first true PCB in 1936 [14]. From simple single-layer boards to today’s complex high-density interconnect (HDI) PCBs with embedded components and 3D architectures, this technology has consistently improved to meet miniaturization and performance demands.

Despite these advances, multilayer PCBs remain accessible only at premium prices, and global supply chain disruptions (such as the COVID-19 pandemic [15] and trade tariffs [16]) have revealed significant vulnerabilities in this approach. And while a quick-turn PCB can be had in as little as a week, the centralized nature of manufacturing creates bottlenecks that can delay development by weeks or months. Decentralized means of PCB fabrication, such as cnc milled and laser ablated FR1 boards, are regularly used by fablabs, research institutions, and companies that need to quickly prototype, but these fabrication approaches usually don’t meet feature parity with ordered PCBs.

1.2.2 Integrated Circuit Packaging

IC packaging has similarly evolved from simple through-hole DIPs in the 1960s to today’s complex system-in-package designs. Surface-mount technology (SMT) became the dominant approach in the 1980s and 1990s, fundamentally changing electronics manufacturing while creating a standardized approach for component integration. And while SMT packages have progressively gotten smaller (BGAs and WLCSPs, 008004 and 006003 [17]; [18]), they’ve fundamentally remained the same as flat devices mounted to a flat substrate.

The biggest change to SMT has been occurring inside the package. Rather than relying solely on transistor scaling, HI focuses on integrating separately manufactured components into higher-level assemblies that provide enhanced functionality and improved operating characteristics. These assemblies are then typically still packaged as SMDs.

1.2.3 Heterogeneous Integration




Figure 1.2: Different Heterogeneous Integration approaches

Heterogeneous integration represents the semiconductor industry’s embrace of modularity principles. By moving away from monolithic designs toward chiplet-based archi-

tectures, manufacturers can mix specialized components fabricated in different process nodes to optimize performance, reduce costs, and improve yields.

The evolution from 2.5D integration (lateral arrangement of chiplets on an interposer) [19]; [20]; [21] to 3D (vertical stacking) [22]; [23]; [24] and now 3.5D [25]; [26] approaches (combining both strategies) demonstrates the industry’s commitment to modularity as a path forward. Companies including AMD, Intel, TSMC, and Broadcom have pioneered commercial implementations of these technologies, with initiatives like Universal Chiplet Interconnect Express (UCIe) and Bunch of Wires (BoW) [27]; [28] attempting to standardize interfaces between components.

However, despite the performance gains made, heterogeneous integration approaches remain largely inaccessible to all but the most well-funded organizations due to several significant constraints.

HI technologies require advanced manufacturing capabilities exclusive to major semiconductor companies and foundries, while the specialized equipment, processes, and expertise involved impose prohibitive costs for smaller organizations. The integration of chiplets demands sophisticated thermal, mechanical, and electrical co-design capabilities that exceed what traditional EDA tools can provide [29]; [30]; [31], and the highly specialized nature of HI manufacturing introduces new supply chain vulnerabilities rather than mitigating existing ones. While proving to be a way forward, this concentration of capability limits innovation to a select few industry giants.

1.3 Supply Chain Challenges and Climate Implications

Electronics supply chain disruptions during COVID-19 created year-long lead times for integrated circuits critical to industries such as automotive and industrial manufacturing [15]. These shortages extended beyond ICs to materials like copper foil [32], impacting PCB production. Global trade tensions and tariffs have further exacerbated these vulnerabilities [16].

Despite government initiatives like the CHIPS Act [33]; [34]; [35] in the United States, robust solutions to supply chain resilience remain largely unimplemented. The centralized nature of electronics manufacturing creates fundamental vulnerabilities that legislation alone cannot address.

Furthermore, as computational demands grow (in particular with the surge in AI, large language models, and training infrastructure), the environmental impact of electronics manufacturing becomes increasingly significant [36]; [37]; [38]. Without more efficient approaches to electronics packaging and assembly, the climate impact of computing will likely worsen dramatically in coming years.

1.4 Accessible Modularity

Previous research in discretely assembled electronics has demonstrated potential for addressing both performance and supply chain challenges. Modular approaches have shown promise for both high-performance computing systems and lower-level electronics packag-

ing, though physical implementations have typically been limited to tens of components rather than the hundreds or thousands needed for practical applications.

The challenge lies not only in technical feasibility but in creating solutions that remain accessible after implementation. Desktop-scale discrete electronics assembly offers a complementary approach to traditional manufacturing; not replacing it, but providing flexibility through localized assembly using standardized components sourced through diversified supply chains.

1.5 Related Work




Figure 1.3: Digital electronic materials history, Neil Gershenfeld

1.5.1 Project Tinkertoy

The concept of modular electronic systems is not entirely new. In the 1950s, when PCBs already existed but most electronics were still manually assembled via labor-intensive methods like wire-wrapping [39]; [40], the U.S. National Bureau of Standards developed Project Tinkertoy [41], a pioneering modular electronics approach that used standardized ceramic wafers as building blocks for electronic circuits. These modules were designed to snap together to form functioning circuits, and although it was eventually superseded by PCBs and transistors, it played a large role in establishing the automation that drives those processes today.


1.5.2 Electronic Digital Materials

In the past two decades or so, more recent work has been done on electronic digital materials, a contemporary rebirth of the ideas that Project Tinkertoy put forth. These implementations have been explored from a variety of angles and abstraction levels by Fredin [42], Langford [43], Ward [44], MacCurdy [45], Hiller [46], and others.

These range from basic conductive and insulating elements, as demonstrated in Ward’s discretely assembled circuits and Langford’s GIK passive components, to component-level systems like BitBlox, to highly integrated programmable processor elements like DICE (Discrete Integrated Circuit Electronics).

These projects have demonstrated the fundamental feasibility of discretely assembled electronics but have generally been limited in assembly scale, automation capabilities, and practical applications. The VMD approach in subsequent chapters builds upon these foundations while addressing their limitations to create a truly scalable volumetric electronics paradigm.

Basic Bulk Materials



(a) Ward’s discretely assembled circuits

(b) Langford’s GIK discretely assembled passives

Figure 1.4: Basic Bulk Materials


Ward [44], Langford [47], Hiller, and others have demonstrated basic electronic digital material systems of different geometries at various levels of automation.

In Jonathon Ward’s geometry [44], he demonstrates conductive and insulating assemblies up to 100 parts, manually assembled.


Additionally, in the process of building an LC resonator, Will Langford assembled an inductor from 520 parts via the time-honored fabrication process of GSWT (graduate students with tweezers, a manual process) [48].

Component Materials

MacCurdy et al. Demonstrated the BitBlox platform with two significant implementations. The first was a two-channel infrared remote control constructed from 17 blocks representing 7 distinct functional types. This assembly was completed using an automated process, though no specific assembly time was documented [45]. In a more complex application, they created a five-channel infrared remote control utilizing 130 blocks



(a) MacCurdy’s BitBlox remote



(b) Langford’s discretely assembled robot

Figure 1.5: Component Materials

across 7 functional types. This larger system was assembled manually and required approximately two hours for both design and assembly tasks [45]. These examples illustrate both the scalability of the BitBlox approach and the potential time efficiency advantages of automated assembly processes compared to manual construction for complex electronic systems. However, the paper also qualitatively suggested that the automated assembly process was only sometimes reliable.

In his PhD thesis, Will also developed a robot assembler for automated assembly, which demonstrates constructing a walking motor from 20 nodes, 6 rigid struts, 4 flexural struts, and 2 actuators; in total, a 32-part walking motor in 8 minutes (4 parts/min, or 240 components per hour (CPH)) [47]. This system could handle 3 different element classes (nodes, struts, and actuators) via toolchanging, and incorporated many quality of life features, such as being able to change part-feeding magazines while jobs were running. However, the assembly rate is slow to build large complex assemblies within a reasonable amount of time; an entry-level desktop pick-and-place (PnP) is 1580 CPH, while this process is 240 CPH.

Discrete Integrated Circuit Electronics (DICE)

The term Discrete Integrated Circuit Electronics (DICE) was initially coined by Will Langford in his PhD thesis [47], and referred to an integrated electronic digital material design that relied on laminate layers to pack multiple conductive layers into an element.

In Tiny-DICE, Zach Fredin designed and fabricated microcontroller assemblies using the smallest COTS components at the time; subcompact mezzanine connectors were used as interconnects for the elements, which were connected in a tetrahedral geometry [42].

His later iteration of Meso-DICE increased the size of the elements to explore more reliable interconnect mechanisms fabricated in-house, and prioritized reliability over feature size.

These geometries were built in 10s of elements, in Tiny-DICE’s case, it was a single type, and in Meso-DICE, it was 2 types.

In general, component and integrated level geometries share similar packaging challenges; to route multiple nets and build-in mechanical properties for alignment and retention, feature overhead must be added for reliable operation.




Figure 1.6: Integrated Materials

Abstraction Levels




Figure 1.7: Component level, in-between basic and integrated

The component abstraction level is the most flexible between basic and integrated, as it can be used for a variety of applications, including computational or power, and provide flexibility in how those applications are implemented. For example, computational architectures could be NMOS, CMOS, diode, etc. Additionally, the component abstraction level can expand to interoperate with basic and integrated abstraction levels easily through a compromise of design requirements; more basic elements can be incorporated by subdividing the component grid and reducing features, and more integrated elements with higher fan-out need only add ports by scaling up in the grid.

Comparatively, operating at the basic abstraction level requires a large amount of assembly to achieve a useful level of complexity; on the opposite end, operating at the processor level greatly reduces design flexibility to only include specific types of computational applications using the same computational architecture, and uses parts with higher overhead that makes it difficult to retool for different use cases.

1.6 Thesis Contributions and Overview

This thesis contributes a volumetrically scalable set of electronic building blocks that can be used to build architecturally agnostic computational logic at the component abstraction level, along with an assembly system for constructing sophisticated electronic circuits. Because the component abstraction level is inherently flexible, this system can be repurposed across multiple abstraction levels, from individual logic gates to integrated structures like operational amplifiers and processor cores.

In Chapter 2, I present a systematic exploration of VMD geometries, analyzing the critical design decisions that govern their effectiveness. I examine manufacturability constraints, electrical and mechanical performance requirements, and feature considerations that influenced the development process. This chapter introduces several candidate geometries, from the initial 4H family to the more refined 4B and 4BI implementations, documenting how each iteration addressed limitations in previous designs. I detail how different tile and connector designs affect mechanical stability, electrical performance, and automated assembly capability, establishing the foundation for a viable VMD ecosystem.

Chapter 3 details the circuit assembly system I developed to construct VMD-based circuits at scale. This chapter examines the motion system architecture, specialized feeders for component storage, substrate design for precise positioning, and end effector configurations optimized for reliable manipulation. I present assembly benchmarks that demonstrate the system’s progression from handling tens of components to reliably placing hundreds of elements with minimal human intervention. This chapter also analyzes potential fault states and error correction strategies critical for scaling to larger assemblies, representing a significant advancement beyond previous work in modular electronics.

Chapter 4 evaluates the performance characteristics of assembled VMD circuits, focusing on two critical aspects: joint reliability and electrical performance. I present analysis of normal force windows and resistance measurements for different contact implementations. The chapter also showcases some example applications through ring oscillator and logic gate implementations. Performance characterization extends to first-principles projections for scaled-down micro-regime implementations, establishing the performance envelope for future VMD development.

Chapter 5 examines the current and upcoming VMD fabrication processes, with particular focus on the development of the 4BIc geometry. This chapter details my process for rapidly prototyping electrical “Connectors in a Day” (eCiD) using commercially available digital fabrication tools, and documenting the challenges and solutions encountered when developing self-aligning connector interfaces that enhance reliability while maintaining accessibility.


Finally, Chapter 6 concludes with a synthesis of the research findings and outlines promising directions for future work, including potential applications in heterogeneous computing systems, superconducting electronics, and scaled-down implementations approaching the semiconductor integration level. This chapter positions VMDs within the broader context of electronics manufacturing evolution and discusses their potential impact on supply chain resilience and sustainable electronics development.

Collectively, these chapters document the development pathway from conceptual design to practical implementation, and aim to establish VMDs as a viable path towards truly volumetric electronics that bridges the gap between conventional packaging technologies and advanced heterogeneous integration techniques.


2 Volume Mount Devices (VMDs)

In the current generation of DICE, there have been several geometries, split into families, explored in the process of discovering a geometry that satisfies fabrication, routing, and reliability requirements.

The current active geometry is 4BI, which is stacked in an alternating fashion. This geometry has been used to assemble an 8-layer, 256 tile mechanical benchmark assembly with only 2 faults that required manual intervention, significantly more elements than previous work at this abstraction level.



(a)



(b)

Figure 2.1: 4BIc0 circuit assemblies: 4BIc0 8-layer mechanical benchmark (254/256 tiles) automatically assembled in ~14min (a); 4BIc0 2-layer stackable full-adder circuit (146 tiles) manually assembled in ~55min (b)

On the way to 4BI, a variety of geometry topologies and configurations were explored in the mm to cm scale (Figure 2.2).

2.1 Geometry Overview

The naming scheme is loosely based on (number of edges)(interconnect)(additional descriptors).

- number of edges: typically 4, though there have been experiments with 3-edge geometries
- interconnect: usually the shape of the interconnect itself; 4H refers to the H-shaped connector, 4B refers to two surface mount contacts pressed against each other, and 4BI represents surface mount contacts and alignment pillars ¹

¹credit to Neil Gershenfeld for creating this naming scheme




Figure 2.2: A variety of geometry iterations




Figure 2.3: Fabricated geometry families

- additional descriptors: shorthand for a feature (i for integrated, p for power, k for kinematic, c for connector, 0 for just the PCB)

In chronological order, the naming definitions of geometries that made it to fabrication (Figure 2.3):

- 4H: 4-sided, H-shaped interconnects
- 4Hp: 4-sided, H-shaped eye-of-the-needle (H-eon) interconnects
- 4B: 4-sided, B-shaped pad-to-pad interfaces
- 4BI: 4-sided, B-shaped pad-to-pad interfaces, with I-shaped alignment pillars
- 4BIc: 4-sided, B-shaped pad-to-pad interfaces, with I-shaped alignment pillars, using compliant connector interfaces
- 4BIc0: a lite-version of 4BIc, using just the PCB. Fabricated and used in the same manner as 4BI.

Table 2.1: Geometry feature comparison

<i>Geometry Name</i>	4H	4Hp	4B	4BI	4BIc0	4BIc
<i>Connection Type</i>	Insertion	Insertion	Pad	Pad	Pad	Pad
<i>Preload Method</i>	Local	Local	Global	Global	Global	Global
<i>Element Types</i>	2	2	1	1	1	1
<i>Dedicated Power</i>	No	Yes	No	No	No	No
<i>Mech. Stability</i>	Low	High	Low	Medium	High	High
<i>Asm. Complexity</i>	Medium	High	Low	Low	Low	Low
<i>Alignment</i>	Connector	Connector	Template	Template	Template	Template
<i>Assembly Force</i>	High	Medium	Low	Low	Low	Low
<i>Joint Reliability</i>	Poor	Good	Poor	Fair	Fair	Good
<i>Scalability</i>	Poor	Fair	Poor	Good	Good	Good

The evolution from 4H to 4BI represented a significant shift in VMD design philosophy (Table 2.1). The 4H family utilized separate tiles and interconnects with through-hole connections requiring local preload, while the 4B family eliminated separate interconnects in favor of surface mount connections with global preload.

The big problem with the initial 4H geometry was a lack of contact compliance and interconnects with non-uniform strength properties. The 4Hp approach offered better joint reliability and uniform strength properties through a single compliant connection design but suffered from high insertion forces and complex assembly requirements.

The transition to 4B significantly simplified assembly by eliminating the separate interconnect elements altogether, but introduced alignment challenges. 4BI addressed these alignment issues through a comprehensive templating system, though at the cost of joint reliability compared to 4Hp.

This evolution demonstrated a trade-off between assembly simplicity and electrical reliability, with each geometry optimized for different aspects of the VMD concept. While 4Hp excelled at creating reliable electrical connections, 4BI proved superior for automated assembly at scale, representing the current active geometry in the VMD ecosystem.

2.2 4Hx Family

2.2.1 4H

4H represents the initial foundation for the current run of DICE geometries, sometimes referred to as “O’s and H’s” (Figure 2.4). Its approach distinctly separates tiles and interconnects as fundamental elements, with tiles designed in square configurations and interconnects adopting H-shaped structures. The system employs two distinct tile categories: functional tiles containing a single component, and routing tiles facilitating net movement along a Cartesian grid.

While the 4H design benefits from conceptual simplicity, it presents several significant limitations. The provision for only a single net per edge makes ubiquitous power delivery particularly challenging; additionally, the system’s requirement for two interconnect types, conductive and insulative, resulted in inconsistent mechanical rigidity, which distorted the lattice structure and compromised assembly consistency.

The H-shaped interconnects lacked geometric compliance, causing all ports to rely solely on press-fit mechanisms and the compliance of bulk materials, which introduced unreliable forces, especially where geometric tolerances were poorly controlled; these tolerance issues proved especially difficult to manage given the large quantity of components needed and the outsourcing of designs to board houses, further complicated by designing at the boundary between reasonable size and achievable tolerance levels. An additional challenge was the single slot intended for two H’s; a single H could leave the system underconstrained and theoretically introduce assembly error.

2.2.2 4Hp

4Hp, previously named o_power, represents an iteration over the 4H design that introduced additional ports for dedicated power delivery and mechanical stability, along with compliant H interconnects (Figure 2.6). This evolution addressed several key limitations of the original 4H approach while maintaining its fundamental architecture.

The 4Hp design achieved notably improved mechanical stability, though the seating issue remained difficult to tune properly. This enhanced stability enabled consolidation to a single H-interconnect type (conductive only), resulting in more reliable mechanical behavior throughout the system. The introduction of dedicated power ports enabled ubiquitous power delivery across the grid, addressing one of 4H’s primary limitations.

By pushing our H-connector element towards more compliant designs, we improved the reliability of each connection and increased the possible lifetime mate cycles. Eye-of-the-needle (EON) compliant designs were heavily used in the connector industry for automotive and other high performance applications. Compared to other compliant designs, EON features were easily scaled down [49]; reducing connector size with respect to tile size enabled us to move away from slots towards holes, increasing connection density on each tile. This allowed us to add dedicated power delivery connections both laterally and vertically, which also worked to improve the overall stiffness and mechanical consistency of the structure during assembly.

These connectors were fabricated from 0.6mm C510 Phosphor Bronze sheet stock, cut






Figure 2.4: 4H dimensions: the tile (a), the H connector (b), and the lattice (c)



(a) 4H tiles



(b) basic 4H led circuit

Figure 2.5: 4H overview


using the Fablight, a 3kW fiber laser (Figure 2.7). Initially, the H-eon connector geometry was compromised and plastically deformed after repeated cycles because the oval was cut-off; as shown in Figure 2.8a, I incorporated several quality-of-life design improvements such that the new geometry had a whole oval that was also above the joining strut, so the pin behaved compliantly. Additionally, the new version had an indent to set the tiebar lower into the body; the leftover stub would protrude from the design and prevent the connector from seating properly, and this feature made the connector behave in a flush manner.

To build basic logic gates and other circuits, we developed a 6-tile library that consists of both functional tiles for components and routing tiles. As shown in Figure 2.9, routing tiles enable connecting nets to power or to selectively direct signals in different directions. This tile system maintains dedicated power ports for each side, located in such a way that structural stability is maintained while being able to selectively populate signal pins.


The introduction of dedicated power ports were not without cost. To avoid adding too many pins to each tile which would make it harder to shrink the design for future, higher density revisions of the geometry, a minimal number of power ports were added, which meant symmetry across the diagonal of the tile, but not axially. This meant that introducing any new tile would also necessitate adding a few variations to enable all possible routing orientations. As a result, additional tiles were introduced for easier implementation of NMOS logic, such as pull-up resistor tiles and FETs covering various degrees of freedom.

Additionally, I designed macro-tiles, which were larger tiles that were ordered in sizes of 4x4, 4x8, and 8x8, used to create sub-assemblies called “blocks”, which enable easier assembly/disassembly, or error correction, for larger structures. The structure is shown in Figure 2.11.


Despite improvements over 4H, 4Hp introduced several new challenges. The dedicated power ports led to significant assembly scale issues, with interconnect pick-and-place op-



(a)




(b)




(c)


Figure 2.6: 4Hp: the tile (a), the connector (b), and the lattice (c)




(a) h-eon-cq




(b) h-eon-fab




(c) h-asm



(d) tile




(e) tiles



(f) 4hp-asm

Figure 2.7: Top: H-eon, from cad to fab, Bottom: 4Hp tile, from cad to fab



(a) left, v1.0; right, v1.1 connectors




Figure 2.9: left, basic 6-tile library; right, implemented power, ground, and signal mapping for tiles




Figure 2.10: 4Hp Batch 009 tiles




Figure 2.11: Hierarchical assembly of tiles and macro-tiles, which enable error-correction at larger, block-sized scales

erations dominating assembly complexity at approximately a 4:1 ratio against tiles. The partially symmetric power ports also led to a combinatorial explosion of additional tiles to maintain design freedom, making designs more complicated and less flexible. Interconnect tolerances proved difficult to control tightly and were challenging to outsource using rapid prototyping fabrication houses. Additionally, interconnect pick-and-place operations proved less reliable than tile placement. The connection mechanism operated contrary to zero insertion force (ZIF) principles, functioning more like “much insertion force” (MIF); this high insertion force would frequently lead to component damage in common misalignment scenarios, as well as loss of stepper motor steps during assembly operations.

2.3 4Bx Family

2.3.1 4B

4B represented a significant departure from the 4Hx family, eliminating the separate interconnect primitive in favor of using tiles only (Figure 2.12). This simplification offered the advantage of requiring only one end effector approach for a single primitive type, streamlining the assembly process. However, the 4B geometry introduced its own set of challenges. Since components remained on the tiles, the assembly process required either inverting the tiles or changing the end effector nozzle size to accommodate varying component dimensions. More critically, this geometry inherently lacked the templating and alignment features necessary for automated assembly, making it impossible to scale assembly without accumulating significant error.




Figure 2.12: 4B design and assembly tests

To optimize the geometry for scaling down to smaller regimes, these new geometries were developed specifically to eliminate interconnect components. While 4B incorporated an initial alignment template, subsequent layers were left under-constrained, inevitably leading to cascading misalignment errors that compromised system integrity and reliability.

2.3.2 4BI




Figure 2.13: 4BI dimensions: 4BI tile (a), 4BI lattice pitch (b), 4BI template top down (c), template corner features (d), and side profile of flexures (e)

4BI represents the current active geometry in the DICE ecosystem, serving as an iteration built upon the 4B foundation, but introducing a template that works for assemblies larger

than a single layer (Figure 2.13). This approach introduces significant improvements in alignment capability, though initial prototypes featured templates lacking compliant alignment features, creating a suboptimal tradeoff between pick-and-place accuracy and constraint. Subsequent iterations integrate flexures which effectively address this limitation. Compared to 4Hp, 4BI significantly improves assembly yield. Like 4B, 4BI’s alternating structure includes gaps that prevent components from interfering with neighboring tiles.

The 4BI approach (Figure 2.14) demonstrates the first version of a geometry that properly utilizes an external structure (the template and the compression cap) for mechanical constraint; this simplifies the pick-and-place process by reducing forces during component placement and relaxing tolerances, as evidenced by later versions of templates with compliant pillars.




Figure 2.14: initial 4BI geometry concept and prototypes (Neil Gershenfeld)

4BI tiles are all 1-layer PCBs, except for special via tiles which are 2-layer. Because dedicated power ports are no longer present in the geometry, tiles avoid the asymmetry problems present in 4Hp, however, the tradeoff is additional tiles required for power delivery within circuits. This leads to a much smaller, manageable set of tiles, which also makes scaling to larger circuit assemblies easier (Figure 2.15, Figure 2.16).

Despite these advances, 4BI introduces several notable constraints. Because the geometry introduces regular pillar features, macro-tiles have less design freedom and face fabrication challenges, since they need to accommodate holes for the pillars. Discrete joint reliability proves reliable for small circuits, but substantially worse than the 4Hp approach as circuits get larger, and the system lacks ubiquitous power distribution ca-




Figure 2.15: 4BI tiles




Figure 2.16: 4BI tiles in gridfinity bins for storage (a), and tiles loaded in a feeder for assembly (b)

pabilities (for now). The 4BI approach relies heavily on global preloading to ensure reliable joints, which significantly restricts structural design freedom; this implementation strongly suggests that gaps can't exist in the lattice structure for reliable global preload to function properly.

In a sense, 4BI is a reshuffling of forces; instead of high insertion force causing troubles during assembly, the same forces must be present post-assembly for the circuit to work.

In the following chapter, I detail the development of an automated assembly system centered around the LumenPnP capable of constructing VMD-based circuits at scale. This system represents a significant advancement beyond the manual or limited automation approaches documented in prior work. By addressing key challenges in motion control, component feeding, substrate design, and manipulation, I demonstrate how VMD elements can be assembled efficiently and reliably. The resulting system not only proves the viability of the VMD concept but establishes a pathway toward scaling to the hundreds or thousands of elements required for practical applications.

3 Circuit Assembly System

To automatically assemble these geometries, a motion system and end effector suitable for reliably manipulating these geometries is necessary. The two most viable approaches are mechanically gripping and retaining by friction, and pneumatic pick up.

In this chapter, the assembly system will be described and specific modifications for each geometry detailed. Geometries will then be compared and contrasted on their ease of scalability.

Scalability refers to ease of automated assembly, including fabrication and assembly of the elements (tiles and connectors) themselves, and assembly of circuits using the elements as feedstock. Yield at each stage is important, as well as cycle times.

To implement applications of sufficient complexity, circuits on the order of at least 100s of elements need to be assembled, as partially demonstrated from the related work described in Chapter 1. In our context, 10s of tiles are enough to demonstrate basic functional circuits, such as a NAND gate or ring oscillator, 100s of tiles are enough to demonstrate a full-adder, an important functional block towards a computer.

3.1 Assembly

Over the course of assembly testing, multiple placement benchmarks were achieved, starting from 4Hp. The majority of the benchmarks were conducted using the 4BI geometry, since it is designed to scale better. These benchmarks are shown in Figure 3.8.


The flow of operations for assembly of VMD circuits is essentially the same as using a traditional pick-and-place, but adjusting the z of successive components for each additional layer (Figure 3.2).

3.1.1 Setup


First, the feeders and substrate are installed into the machine. These are currently Form 4 Black V5 resin prints that have been printed flat on the bed; this makes sure that surfaces are as flat as possible from the printer and precludes need for support. They are installed using M3 screws and nuts.

Then, tiles are loaded into feeders, organized by type, with consistent orientation.

Next, the substrate (build platform) is installed. I try to keep the feeders and substrate relatively close to reduce travel time. If the cameras are being used, that's a good place to locate them near, but the stock configuration seems to put more tension on the camera/main platform for the LumenPnP so I prefer to use the auxiliary platforms for my substrate and feeders.




(a) 4Hp 8x8 template, 14/16 tiles



(b) 4BI 8x8x2 template, 41/43 tiles



(c) 4BI 8x8x5 template, 157/160 tiles



(d) 4BIc0 8-layer template, 254/256 tiles

Figure 3.1: Benchmark tests




Figure 3.2: The assembler picking and placing from 4BIc0 feeders

With everything installed, I pass a pnp job (.pos, or .board.xml, the native OpenPnP format for jobs) into OpenPnP, run calibration on the feeders and substrate, and press go.

3.1.2 Operation

The head moves to the feeders, picks up a part, goes to computer vision, checks part orientation, goes to substrate, places part, and repeats this process until job finish, or as a fault occurs. The basic steps are shown in Figure 3.3.




Figure 3.3: Basic PnP operation (without computer vision): picking (a), moving (b), placing (c)

3.2 Motion System

Off-the-shelf open source desktop pick and place machines have matured to the point that it makes sense to use that as a starting point and see how far assembly can be scaled using this motion system. Designing around an existing system allows us to allocate more resources towards geometry development relative to the assembly system.

The LumenPnP (Figure 3.4) is a CoreXY motion system with a relatively low z travel, optimized for 2D pick and place operations. In our packaging system, the LumenPnP is used for 2.5D assembly.

3.3 Feeders

Feeding is an integral part of the assembly process; as assembly complexity increases, the need to feed higher volumes of primitives increases as well. Resetting jobs manually with tweezers becomes increasingly limiting without automation in the loop.

3.3.1 Passive Tray Feeders

Passive tray feeders were found to be the best fit for initial pick and place operations, initially designed using solidworks, then formalized with feeders.py (cadquery), a parametric script subsequently used to generate arbitrary feeders for both the pnp process and other fab work holding and feeding.

The geometry and its associated infrastructure (feeders for holding the parts, end effector to pick and place the parts, template to receive the parts), were initially designed





Figure 3.4: LumenPnP

in SolidWorks, a parametric point-and-click MCAD tool. In some cases, such as the 4Bx family templates and feeders, designs were then migrated to CadQuery 2, a parametric scripted MCAD tool. CadQuery 2 offers performance and organizational advantages, namely the python support which makes headless automation and scripting with external libraries easy (versus something like the SolidWorks or Fusion360 API). This supports our bottom-up (`dice_cad`) and top-down design tools designed by my colleague Erik Strand. Feeders are shown in Figure 3.5.


3.3.2 Adding Depth

As assembly transitioned from 10s to 100s, depth was added to feeders to save on xy travel as well as pnp work area. For 4BI, a depth of 3 tiles (of 1.6mm each) was found to be a reasonable capacity; larger than 3x became difficult to load tiles in the correct orientation.


Packaging changes introduce additional challenges for feeding; simple 2d tiles (such as PCBs) have a predictable z-thickness, but packages with compliant flexures have variable z-thickness under load (4BIk) (such as during pick and place) and packages that have been balled (4BIb) (BGA) are subject to additional tolerance stackup based on control of solder sphere tolerances. For exceptionally varied packages, tile planarity also becomes an issue (4BIb?), which needs to be controlled by the feeder geometry itself.



(a)



(b)



(c)



(d)

Figure 3.5: Parametric feeder design in CadQuery: 30 element capacity 3x10_feeder (a), 60 element capacity 3x20_feeder (b), 60 element capacity 6x10_feeder (c); evolution of feeder concepts for 4Hp, from trays to magazine dispensers (d)

3.3.3 Alternative Feeders

There are many other types of feeders beyond passive tray feeders, such as tape feeders, or vibratory bowl feeders. [50] Traditionally, 1000s of components are likely to be packaged in tape-and-reel. To maintain compatibility between our tiles and industry standard tape-and-reel for feeding, 8mm and 12mm wide tapes would be ideal. Tape and reel introduces additional packaging challenges and tolerance stackup. An alternative approach to active feeding of arbitrary geometries is vibratory feeders. Vibratory feeders come in a couple of flavors, including dish and linear. The main advantage of a vibratory feeder setup is being able to sort loose parts for predictable pick and place. These systems can be combined with computer vision for additional or complementary error correction of part orientation.

3.3.4 Resetting Jobs

The current tray feeders work reliably and have room for increased density. However, while they are easy to populate for jobs of 10-100 elements, they quickly become tedious to populate with tweezers, both the conventional and vacuum types. This is especially true when considering orientation (rotationally and flipped) matters. Populating 280 tiles into 7 feeders (4x10) took ~40min by hand, and it happened that this job was purely a mechanical benchmark with dummy tiles that didn't care about orientation; properly orienting tiles would add additional time.

One possible way to tackle this problem is to observe that feeders and templates are both forms of tile storage. Since a feeder can be used to feed tiles for a template, why can't a template be used to feed a feeder? Feeders are optimized for packing tiles in a non-functional state for the purpose of having tiles removed, while templates are optimized for constrained structural lattices in a functional state for retaining tiles.


Storage solutions for these opposite states are easy to design for separately, but difficult to bring together; while removing tiles from a feeder should be low-force, removing tiles from a template should be difficult.

3.4 Substrates


Substrates are large feedstock used as a platform for mechanically constraining and sometimes electrically routing elements assembled into circuits. Related work from Langford [43] and [46] demonstrate uses of substrates for preloading and alignment of elements as well. Two major types of substrates were explored over the course of the geometry design study, the macro-tile from 4Hp, and the template from 4BI.

3.4.1 4Hp Macro-tile


The 4Hp macro-tile (Figure 3.6) both provided mechanical support and enabled electrical routing and made breaking out signals easy. Additionally, it was designed with edge tolerances in mind such that tessellating macro-tiles was possible, enabling larger circuits beyond a single macro-tile. However, it made automation difficult; the through-hole



(a) 4Hp 8x8 template, 14/16 tiles



(b)



(c)


Figure 3.6: 4Hp macro tiles as substrates

connections used to interface with H-eon connectors required higher tolerances, which weren't possible to achieve at scale through the fabrication of H-eon connectors. In addition, 4Hp relied on high insertion forces; this caused the assembler to skip steps during insert, which meant rehoming between each place. In cases of misalignment, this also caused the macro-tile to kinetically discharge chunks of FR4.


Because the macro-tile was a PCB following standard PCB design rules, it wasn't possible to improve lead-in by chamfering the through holes.

3.4.2 4BI Template


In contrast, the 4BI template is effectively a zero insertion force system. Some common configurations are shown in Figure 3.7.



(a)



(b)




(c)


Figure 3.7: Parametric template design in Cadquery: 9x9x10 template, 320 element capacity (a), 9x9x3 template, 96 element capacity (b), 17x17x10 template, 1280 element capacity (c)

Like their counterpart feeders, 4BI templates were also initially designed using solidworks, then formalized with 4BI.py (cadquery). The current tradeoff between 4BI templates and 4Hp macro-tiles are that while the templates are easier to parametrically generate using 4BI.py, 4Hp macro-tiles are better at tessellation with standardized sizes. Adding this capability to 4BI is on the roadmap.


Initial templates used rigid pillars for locating and constraining tiles in xy, subsequent templates implemented compliant pillars that were more forgiving of coarser tolerances, and were important for maintaining pad to pad alignment.



(a)



(b)



(c)

Figure 3.8: Templates in use for various tasks: early rigid template with poor tolerance (a), high-temp reflow variation from soldering tests (b), matured design using Form 4 Black V5 (c)

Over the course of these two jobs we identified problems as well as solutions to increase reliability necessary for automatic placement and avoiding the need for manual intervention. Initial templates featured rigid, non-compliant pillars that ran into tolerance issues; too tight and tiles would bind on the pillars, preventing proper seating and removal, too loose and neighboring pads would misalign and miss, resulting in open circuits. Subsequent templates adopted compliant pillars, which enable tighter tolerances. These new templates guarantee alignment and help retain tiles after placement, which

makes assemblies less fragile.

4BI templates also require caps for compression to both preload the entire circuit and increase surface area of contact interface. This concept is similarly applied by [43] in his testing of active components in section 5.2.

3.4.3 Material Choice

During iteration, flexible, elastic resins were considered (such as Form 4 Flex 80A). While these materials offered excellent compliance, they unfortunately exhibited unwanted tackiness that interfered with assembly operations. In addition, the slow springback after deformation was too slow for reliably maintaining xy alignment.

After evaluation, we ultimately stuck with rigid resins with integrated flexures instead of fully flexible templates. This alternative provided more responsive mechanical behavior with less surface adhesion, without compromising the structural rigidity needed for effectively constraining the system during assembly.

Another variation was the Form 4 High-Temp resin. This variation was used for solder experiments, where 4BI tiles were loaded with solder prior to assembly, then subject to compression and heat to initiate reflow. For implementing this approach, High-Temp resin was required, even for low-temp BiSn solders @ 138C.

3.5 End Effector

End effector went through multiple design iterations, but ultimately with both element consolidation and proven effectiveness from stock solutions, the standard vacuum end effector with stock N24 nozzles has persisted up to 4BI.

3.5.1 Claw Gripper

For 4Hp, a claw gripper design based on Zach Fredin’s design for Meso-DICE was implemented, targetting tiles. This design uses a cam driven by a servo to guide the fingers of the gripper in and out of compression with a tile.

These systems both utilize a bottom camera for registering tiles to the global grid and a load cell for verifying successful placement of elements.

Although the vacuum pick-up strategy is sufficient for adding tiles to a structure, it doesn’t work well for removing tiles from a structure. We have created a prototype claw end effector, which aims to effectively both add and remove tiles, shown in Figure 3.9. Future work will extend the automated assembly to 3D, and begin increasing the component count.

Several gripper iterations targeting the 4Hp geometry were fabricated, but due to binding issues and effectiveness of the vacuum nozzle approach, the claw has been parked for now.




Figure 3.9: Evolution of tile gripper designs for 4Hp

3.5.1.1 Passive H-gripper

An early 4Hp H-eon concept used an electromagnet actuation strategy to clamp and unclamp elements was aimed at the H-eon connector. Later, this approach was simplified to a passive gripper, matching the profile of the slots in the tiles themselves and relying on the compliance of the H-eon connector itself for pick-up, shown in Figure 3.10.

Since moving to 4BI and beyond, there hasn't been a need for the H-eon end effector.

3.5.2 Vacuum Nozzle

Yohan Guilamard worked on custom multi-orifice nozzles targeting corner pads on 4BI.

3.5.2.1 Registration and Alignment Strategies

Our initial design assumed the nozzle would deflect during approach, providing a form of elastic registration. However, the final solution involved designing the template to deflect through carefully engineered flexures, which significantly reduced misalignment issues during the assembly process.

3.5.2.2 Nozzle Material

Stock nozzles are made out of steel, which are physically robust to shocks, which can be somewhat common in the assembly iteration cycle. This is especially true of the thin walls and small channels. However, stock nozzles don't assume high anisotropic placement (deep placements), so combining stock nozzles with 4BI templates leads to compromises (with 1.6mm tiles, this is a limitation of 5-7 layers, depending on template alignment feature overhead, such as chamfers for the first 2 layers).

More durable resins may also improve performance and make this feasible; Black V5's strength properties are insufficient for the job, they are too brittle. Flex 80A deforms elastically rather than breaking, but we found this material too sticky, which is problematic for consistent pick and place (emphasis on the place).





Figure 3.10: Passive gripper end effector for 4Hp H-eon connector

3.5.3 Sealing the Vacuum

We identified an important assumption mismatch regarding tile materials. Milled FR-1 tiles typically lack soldermask, which can lead to gaps between copper traces and milled channels, potentially compromising vacuum pick-and-place efficiency. Despite this theoretical concern, the LumenPnP system demonstrated consistent performance between both FR-1 and PCBWAY FR-4 tiles (with soldermask). However, other suction systems experienced reliability issues with pick-and-place operations, either due to nozzle geometry (not completely flat, or conformal) or due to reduced air flow from less powerful pumps.

3.5.4 Nozzle Geometry

Due to unwanted tilting with the stock nozzle, time was spent looking into optimizing the nozzle geometry to enforce planarity of the tile.

Optimal nozzle geometry spans most of the tile surface area, with contact points strategically positioned at corners to ensure even distribution of forces and promote planarity of the tile during handling. Supporting complex geometries can introduce asymmetry, especially when incorporating fine features such as small channels and thin walls—both common elements in vacuum pick-and-place nozzles.

During experimental comparisons between stock nozzles and custom-designed alternatives, we observed that tile stability remained good as long as the resulting contact plane (comprised of single or multiple orifices) covered a majority of the tile’s surface area. For stock nozzles and 4BI geometry, N24 nozzles are the best fit.

A possible iteration that would remove the limitation within this system (7+ layers) would be to make a custom nozzle that extends the geometry of the stock N24 nozzle, using a metal fab process (possibly outsourced, likely machined by eg PCBWAY CNC).

3.5.5 Tile Surface Planarity and Area

Additionally, elements explored in this thesis have all been components on top of the tile. This makes a big difference in contact surface area for the vacuum end effector; picking the bare PCB maximizes surface area and enables us to use the N24 nozzle, while picking the side with a component either forces us to use a much smaller nozzle (N45) or use a much larger nozzle to fit the component inside the nozzle (N08), which ironically has restrictions fitting inside the template itself.

In a future iteration, integrating the functionality into the tile would offer more flexibility (much like an IC that behaves in a VMD, rather than SMD, manner).

3.6 Error Correction

Statistically speaking, in the case of scaling from 10s to 1000s of parts, it’s not a question of if defects will occur, but rather how many defects will occur. Robust error correction mechanisms are required to properly assemble circuits that function at scale. There are multiple strategies to address these issues: error correction during assembly, error

correction after assembly, and ease of assembly/disassembly to such that these solutions are practical.

During assembly, it has been observed that tiles can be placed incorrectly, resulting in a few different error states, such as the ones in Figure 3.11.




Figure 3.11: 4BI error states: failure to rotate, tile never enters template (a), tile not completely flush; opportunity to correct but may interfere with adjacent placements (b), tile placed on its side; tile enters template but is invalid and affects adjacent tiles (c)

These error states range from inconvenient to catastrophic; some errors demonstrate an ability to recover while others prevent neighboring tiles from achieving a valid state. These errors can be largely mitigated by robust error correction mechanisms during assembly. The current approach heavily relies on computer vision for orientation and center detection of tiles in order to pickup from the feeder and correctly place into the

template. This resolves to properly recognizing the boundaries of the tile which is usually based on sufficient contrast of identifying features; in this case reflection of conductors facing the camera. For purely mechanical tiles, reflection of the bulk substrate is the identifying feature (eg white resin performs while black resin lacks sufficient contrast). By improving estimation of these features and making sure angle corrections are carried out within tolerance, many of these scenarios can be avoided.

Additionally, end effector geometry plays a significant role in distributing forces evenly during pick-and-place operations. For instance, using a larger N24-sized nozzle covers $\sim 70\%$ of a 4BI tile's surface, providing more stable support compared to an N08-sized nozzle that only covers $\sim 25\%$. When the tile is supported at a single central point, such as with a small nozzle, it is more susceptible to tilting or rotating due to uneven force distribution and higher moments at the edges, especially if a collision occurs with an alignment pillar. An example of a tile tilting due to uneven solder and small contact surface is shown in Figure 3.12a. In contrast, supporting the tile at multiple points, like its four corners, reduces these moments and helps maintain stability when placing the tile into a container, as shown in Figure 3.12b. Uneven soldering also introduces errors exacerbated by central support versus edge support.




Figure 3.12: Vacuum end effector: Tile tilting due to uneven solder, center support (a), CAD of end effector supporting tile on edges, even force distribution (b)

A 2-fold approach for coarse and fine positioning of geometry is useful for reducing error correction defects and improving scalability; for example, converting the pillars in the template to improve compliance reduces tolerance requirements for coarse positioning, while the pillars themselves are used for fine positioning.

Post assembly, it has been observed that poor tolerance stackup contributes to accumulated error, eventually leading to intermittent contacts. One possible solution is following a physically hierarchical approach; individual templates are parametric and can be sized such that a limited number of tiles are placed (eg 4x4x4 or 8x8x8), and finished block assemblies then assembled together to form a much larger assembly. The faces of the block assemblies operate as interposers that then pass signals and power across blocks; the interconnect system here can be designed for quick and robust connection/disconnection to troubleshoot and correct malfunctioning blocks.

3.6.1 Overconstraint and Compliance

Part alignment is key towards functioning circuits, local (part-to-part) and global (template-to-part) alignment have been the two main strategies. Deciding between the two strategies impacts tolerance stackup. If tolerance stackup is not properly managed, physical defects can occur due to misalignment and undermine mechanical and electrical performance. An example from Tiny-DICE is in Figure 3.13.




Figure 3.13: SlimStack connectors from Tiny-DICE, broken from overconstraint [42]

3.7 Comparative Assembly Rates




Figure 3.14: Will Langford's scaling throughput graph [47]

Will Langford's analyzed relative rates of assembly between natural and artificial means (Figure 3.14), showing the capability gap that needed to be crossed to reach parity with

resolution vs number of elements assembly [47].

Although we are not yet ready to address the massively parallel assembly implications from Will’s thesis, we can still derive reference points for contextualizing our own assembly throughput; in his chart, commercial pick and place machines are projected to be among some of the highest throughput artificial means of assembly. Although all referred to as pick and place machines, there is a wide range of capability, covering different regimes.

Table 3.1: Pick and Place Machine CPH Comparison spanning multiple categories

<i>Machine</i>	Claimed CPH	Technology / Notes
<i>Hand Assembly</i> ¹	~100	Manual placement by human operators.
<i>LumenPnP</i> [51]	1,580	Open-source, DIY-friendly desktop PnP.
<i>Neoden 3V</i> [52]	~5,000	Budget small-batch machine.
<i>Juki RS-1R</i> [53]	~42,000	High-speed modular industrial machine.
<i>Fuji NXT III</i> [54]	35,000–37,500	High-end modular system; each module rated.
<i>Yamaha YSM40R</i> [55]	~200,000	High-speed modular rotary head system. Discontinued.
<i>ASM Siplace SX</i> [56]	~80,000 per gantry	Industrial, configurable for very high speeds.

From Table 3.1, there are a few rates that carry special attention. For most of the related work section discussed in Chapter 1, and most of the iterative work required to build to this thesis, Hand Placement was the primary driver for circuit assembly. Even as assembly improves, augmenting Hand Placement for semi-auto rather than full-auto solutions may be beneficial for iteration, especially if tile geometry changes significantly and requires new end-effectors. Often, effective hand tools translate to effective end effectors, and vice versa.

From here, we move on to the LumenPnP itself. The current actual PnP speed observed from various PnP jobs is about 2/3rds of the advertised max; the LumenPnP head itself is capable of 2 nozzles instead of just 1. The typical end effector configuration uses nozzles of two different sizes, which is likely used to achieve the advertised max for standard PCBs. In our case, we may need to install two N24 nozzles to achieve the advertised max for assembling circuit volumes.

Finally, the Yamaha YSM40R serves as a theoretical conventional limit for PnP throughput. This highly specialized PnP is no longer sold, but holds the record for fastest PnP that has been commercially available. More modular models have since taken its

¹Note that it was difficult finding a source for manual placement CPH; instead, 100 CPH is conservatively estimated from my manual placement job for the 4BIC0 full-adder, which was ~150 tiles in ~55 min.

place, though they have a lower throughput, they are more flexible for reconfiguring to different types of jobs.

These rates of interest have been consolidated in Table 3.2.

Table 3.2: Selected PnP machines for projections

Description	CPH (Components per Hour)
Hand Placement	~100
Current Actual PnP Speed	~1,000
LumenPnP Advertised Max	~1,500
Yamaha YSM40R (Discontinued SOTA)	200,000

Using these selected PnP rates, a projection tradeoff can be made, optimizing for increased serial speed, or increasing the number of machines in parallel, or both. Projections are shown in Table 3.3.

Table 3.3: Parallelization pnp projections

Scale	1x Machine	10x	100x	1,000x	10,000x
1x current	1,000	10,000	100,000	1,000,000	10,000,000
2x current	<i>2,000</i>	20,000	<i>200,000</i>	2,000,000	20,000,000
5x current	5,000	50,000	500,000	5,000,000	50,000,000
10x current	10,000	<i>100,000</i>	1,000,000	10,000,000	100,000,000

3.7.1 Possible Parallelization Pathways

Increasing assembly throughput can be accomplished through several distinct approaches. One method is to increase serial speed for a single assembler, which can be achieved by transitioning to a stiffer motion system capable of handling faster speeds without sacrificing resolution. Low-backlash systems could enable this improvement; the effectiveness is heavily dependent on the specific motion system and hardware implementation. Fundamentally, there is a tradeoff between speed and precision, which sets a hard limit on how fast a machine can usefully go.

Another approach involves implementing simultaneous parallel pickup from a single end effector. This strategy could dramatically increase throughput; however, it presents a challenging control problem to ensure reliable selective pickup for each element being handled. Additionally, this method may introduce more motion overhead to coordinate elements into the correct location for the assembly head.

The VoxJet system exemplifies an innovative parallel assembly approach. Developed by Hod Lipson and Jonathan Hiller, this VoxJet precursor utilized acrylic non-conductive (and conductive, though not emphasized in this context) spheres as building elements with inherent self-aligning properties. The system employed an electrostatic end effector capable of manipulating multiple elements in parallel, reportedly handling thousands of insulating spheres simultaneously [57] [46]. There has yet to be a way to apply a similar

approach to integrated devices for 3d volumes, which is what this VMD approach aims to satisfy.

Finally, reducing the cost of individual assemblers and deploying multiple systems operating in parallel can effectively balance the performance limitations of any single machine. This approach offers additional benefits including increased feedstock flexibility and enhanced redundancy in case of assembler failure. However, it introduces a different category of challenges, shifting focus from coordination within a single machine to sophisticated orchestration between multiple assembly systems.

In the next chapter, we evaluate the joint and circuit performance of these geometries. Additionally, projections are made for scaling the feature size down, which will be necessary for increasing achievable VMD performance.

4 Testing and Performance Overhead

In this chapter, we investigate joint reliability and circuit performance.

Joint reliability mostly refers to resistance of joints, as well as the Normal Force Window, which is the range of preload forces necessary for the connection to behave electrically reliably. May involve cycling tests.

Circuit performance refers to the parasitics of the assembled circuit as compared to using conventional means such as monolithic PCBs or breadboards.

Additionally, analytical performance projections are made showing smaller tiles (micrometers, nanometers) will likely lead to better performance.

4.1 Joint Reliability

4.1.1 4Hp H-eon Elements

Because 4Hp used compliant contacts in the form of the H-eon elements, it had functionally reliable contacts. During initial prototyping of H-eon elements, many parametric sweeps were made on insertion force, reusability, and insertion depth. These traits were balanced with ability to fabricate.

A key problem that led to 1-use, or sometimes simply unreliable contacts were geometries that appeared to be compliant, but instead deformed plastically after the first insertion. It was later identified that the geometry itself compromised the design; the compliant eye-of-the-needle feature was only partially formed. An iteration was created (v1.1), allowing more clearance for the feature and enabling reusability.


Unfortunately, the insertion force and connector cycles were not rigorously characterized before we switched to 4BI for its superior automatic assembly characteristics.

4.1.2 4BI Pad-to-Pad


4BI connections are formed by compressing two pads from adjacent tiles together. This is facilitated by global compression of the assembly, as described in the previous chapter. Compared to 4Hp, this approach is much easier to pick and place leading to significantly improved assembly reliability, but has been a source of intermittency for developing circuits.

4.1.3 4BI Pad Geometry


Initial experiments used rigid pillar templates, which were more tolerance sensitive, and frequently led to misalignment of neighboring pads. The geometry of the pads themselves contributed to this problem.



(a)




(b)



(c)



(d)



(e)

Figure 4.1: 4BI experiments to try and solve intermittency: Spring-loaded compression cap (a), solder and reflow (b,c), experimenting with plating finishes (d), introducing contact geometry via leadframes (e), mechanical decoupling of cable strain

Pads were 1mm x 1mm, and increased to 1.5mm x 1.5mm for improving contact surface area (Figure 4.3), as well as enabling solder reflow tests without modifying template pitch, good for backwards compatibility with existing parts. 4BI geometry follows a 5.655mm pitch, as shown in Figure 4.2.




Figure 4.2: 4BI pitch





Figure 4.3: At 5.655mm pitch, 1mm x 1mm pads (4BI-b001, b002) could easily result in only ~25% pad overlap with unconstrained rotational error. 1.5mm x 1.5mm pads (4BI-b003+) increases this to ~50%, with little downside. Reducing the pitch would improve this even more

4BIc, a new variation on 4BI (the c stands for connector), was then introduced, incorporating contact geometry via leadframes. This is to reintroduce compliant contacts

back into the main branch, which should be more shock and vibe tolerant than bare pads. The manner in which they interface can be seen in Figure 4.4. Unfortunately these leadframes are still experimental, and not yet complete.



(a) the 4BIc “leaf spring” geometry



(b) 1.5mm orthogonal pads (4BIc-b001)

Figure 4.4: 4BIc contact improvements

Despite lacking leadframes, the bare tiles have still been improved. Designs using the bare tiles are named 4BIc0, to indicate a variation of 4BIc without full features. To facilitate easier leadframe design, the pads were rotated such that adjacent contacts are no longer aligned by their diagonals, and instead orthogonal to each other, improving stability when neighboring tiles aren’t fully populated.

4.1.3.1 Resistance Testing

In 4BI-b004, a series of tiles were ordered with different plating finishes, which can be seen in Figure 4.1d. This was motivated by the idea that different material finishes with softer mechanical properties would benefit electrical performance; softer metals would deform easily and have higher contact conformity, leading to lower contact resistance. Additional connector-oriented features also influenced choice of plating finishes.

HASL is generally the go to for economic PCBs, and has been the default finish for test elements. HASL lead and unleaded perform in the same price range and performance domain, so only HASL lead was used. ENIG and ENEPIG were selected for their electrical properties, resilience to oxidation, and wear characteristics. Copper was quickly excluded due to quick oxidation. Phosphor Bronze (C510) is a copper alloy typically used for connector contacts requiring spring properties, great for compliant contact designs. It isn’t readily available as a plating finish from PCB houses like PCBWAY, but easily available as sheet stock from vendors such as McMaster-Carr.

A test setup was then devised to test individual tile pad-on-pad performance. A 3d printed jig (Form 4 Rigid 4k) was used to hold a tile each, and wired to a 4-wire resistance testing setup. This design under test (DUT) would then be subject to an instron, used to compress two adjacent pads from each tile together. Because the instron uses an offline computer, system clocks were used to synchronize resistance and load measurements, accurate to a second.

Each jig holding the tile is loosely constrained in x and y using alignment features, also commonly used for mold making and other tooling registration steps, such as the ones

used to fabricate 4BIc. This was designed this way to avoid potential binding or other overconstraint conditions that would compromise the experiment, ensuring only normal forces between the two pads are measured.

Pyrallux copper tape was used to add well-constrained, solderable conductive paths to break out the tile connections to decouple cable strain from the tile, which was fixed in place using UV glue. The test jigs can be seen in Figure 4.5.






Figure 4.5: Test jigs for single tile pad-to-pad resistance tests: CAD model of test jig (a), soldering stranded wire to pyralux, solder joint to tile (b), test jigs fabricated for each of the 4 test finishes (c), test jigs in their closed state, ready for testing (d).

To install the jigs into the instron, double-sided nitto tape was used to affix the bottom and top halves of the jig to the base and loadcell. It was necessary to affix the bottom to a piece of flat mic6 stock, since the wire strain from the kelvin probes would tilt the setup, oriented the wrong way. The mic6 stock was left loose, to underconstrain the system so the registration features can be aligned. This can be seen in Figure 4.6.

The breakout assembly (kelvin probe $\rightarrow$ wire $\rightarrow$ solder $\rightarrow$ copper tape $\rightarrow$ solder $\rightarrow$ tile $\rightarrow$ tile on-board resistance) series resistance was measured using a 4-wire setup prior to the experiments, such that parasitic resistances are accounted for while measuring resistance of the pad-to-pad joint itself. This parasitic can be seen in the test setup schematic, shown in Figure 4.7. This parasitic resistance measurement was conducted on each of the two halves of the jig (top, bottom), and there were 4, for a total of 8 separate measurements.



(a) test jigs taped to the load cell and platform (b) test jig under compression, up to 100N

Figure 4.6: pad-to-pad resistance vs compression force test setup





Figure 4.7: 4-wire test setup, measuring $R_{PadJunction}$


Table 4.1: Parasitic resistance measurements for each sample

Sample	Parasitic Resistance Sum (Ω)	Parasitic Resistance Sum (m Ω)
1	0.01898 Ω	~18.98 m Ω
2	0.02345 Ω	~23.45 m Ω
3	0.03004 Ω	~30.04 m Ω
4	0.02669 Ω	~26.69 m Ω

Then, each sample was subject to testing, which displaced at a rate of 2mm/min to the end condition of 100N compressive force. Data can be seen in Figure 4.8.



(a) raw measured data



(b) parasitics subtracted

Figure 4.8: Plating finish test results, resistance vs compression load

Subtracting the parasitics netted negative values, suggesting that the resistances in the system for individual joints are below the noise floor of my measuring equipment. This makes comparing finishes purely on resistance impossible, except that all of the finishes appear to perform adequately. However, the data does provide insight into activation forces. The Normal Force Window of all 4 samples was $>50\text{N}$, with an intermittent but still acceptable period from 40-48N. This can be seen in Table 4.2.

Table 4.2: the resulting Normal Force Window for all 4 samples

Sample	Min Resistance ($\text{m}\Omega$)	Stable Force Range (
2b-hasl-lead	16.01	50.15 N $\rightarrow$ 101.50 N
2b-enig-1u	15.23	48.72 N $\rightarrow$ 100.55 N
2b-enig-2u	17.11	48.60 N $\rightarrow$ 101.03 N
2b-enepig	14.34	49.32 N $\rightarrow$ 100.07 N

4.1.3.2 Cycling

For testing reusability, a cycling test is necessary to see how many cycles a connector contact can undergo before behaving differently.

A preliminary cycling test on the HASL-lead sample was conducted to evaluate joint reusability. However, the test inadvertently operated under a strain-controlled mode rather than force- or displacement-controlled cycling, resulting in peak compressive loads exceeding 400 N and triggering the 500 N cutoff condition, which caused the sample to fail after four cycles. While this result highlights the sensitivity of the joint to over-compression and underscores the need for controlled force cycling in future tests, it does not reflect performance under nominal operating conditions. Future work will focus on properly constrained tests required to determine realistic reuse lifetimes.

For a more realistic scenario, a continuity test was implemented using 4BI-b004-2b tiles, to create a “snaking” trace across one of the 8x8 templates, as seen in Figure 4.9. In total, 14 joints were used in the continuity test. Due to time, this test was not conducted and will also be left for future work.




Figure 4.9: Continuity test setup: circuit assembly (a), setup (b), circuit in dice-cad (Erik Strand) (c)

4.2 Circuit Performance

Multiple circuits were built over the course of development. These spanned from basic benchmarking tests, such as continuity tests, then simple operational circuits such as ring oscillators, and finally logical circuits like nand gates and adders, with the intent to eventually build an entire RISC-V CPU.

4.2.1 Ring Oscillators

The ring oscillator has almost been the functional hello-world of each geometry. It demonstrates basic conductivity, some functional components, and a quick way to benchmark relative parasitics through the resonant frequency of the circuit. It is a circuit made of an odd number of inverters connected in a loop, where the signal continuously toggles due to propagation delays, producing a periodic square wave.

A 4Hp implementation can be seen in Figure 4.10. An equivalent ring oscillator design has been made for 4B1c0 (Figure 4.10d), but it has yet to be tested.




Figure 4.10: 2MHz resonant frequency 4Hp ring oscillator circuit: ring oscillator schematic (a), 4Hp ring oscillator circuit (b), oscilloscope shot of circuit ringing (c), 4B1c0 ring oscillator design in dice-cad (d)

4.2.2 Logic Gates

Basic logic gates, as shown in Figure 4.11, can be put together to build many other complex functional blocks, such as half-adders and full-adders, to an ALU, and ultimately culminating in a computer. NANDs (or ANDs in diode logic) are a very commonly used logic gate. We use NANDs and NORs in our more complex circuits.

4.2.3 Full Adders

A 4Hp full-adder circuit (Figure 4.12a, Figure 4.12b) was manually constructed by Shravika Pendyala, which itself is composed of two half-adders assembled on 8x8 macro tiles. This represents a more complex logic application circuit.

In 4B1c0, I designed a half adder in dice-cad (Figure 4.12c). However, given the comparable tile counts, Erik and I opted to carry out implementation of a full adder instead (Figure 4.12d). A stackable full adder (Figure 4.13) was designed by Erik Strand using dice-cad, which I then modified by breaking out connections and fitting it to the 16x16 template. Additional filler tiles were added to improve mechanical stability. Setup







Figure 4.11: 4Hp logic implementations: 4Hp diode AND gate (a), diode AND gate schematic (b), 4Hp nmos NAND gate (c), nmos NAND gate schematic (d)




(a)



(b)




(c)




(d)

Figure 4.12: Adder circuits: Full-adder schematic (a), 4Hp Full-adder implementation (Shravika Pendyala) (b), 4BIC0 half-adder (c), 4BIC0 full-adder (d)

was done to make the pattern pick and placeable, but due to time, I manually assembled the circuit shown here. This took about ~55 minutes, and was done using a straight tweezer.



(a)



(b)

Figure 4.13: ~146 tile 4B1c0 full-adder circuit assembled (a), and in test-setup (b)

For preliminary tests, the power and ground nets were measured for continuity and

resistance; the Normal Force Window required for activation was $>100\text{N}$. The ground and power net measured $\sim 0.1\text{mOhms}$. Unfortunately, the center circuitry appears to be intermittent, and requires further work to diagnose. There are multiple potential issues preventing the circuitry from working reliably, one possibility is difficulty guaranteeing even force distribution due to warping of the resin template and cap, or non-planarity from the tiles and their stackup, or both.

Discretizing the cap itself and redesigning it for additional compliance (or adding a compressible layer) are potential solutions to the above problems, but the lack of compliance from the tiles themselves may make required forces difficult. Even after applying some of these solutions and increasing to 300N didn't show improvement.

This is likely where 4BIc could help the most; compliant leaf-spring contacts on each joint interface may help evenly distribute forces while forcing contact at forces well below bare 4BI contact pads, further investigation and process development here is required to scale functional circuits. In the next chapter, I will talk about the manufacturing process for 4BIc, which further builds upon 4BI's ease of assembly with ease of conductivity.

4.3 Scaling Performance Projections

As our geometry shrinks from the millimeter to the micrometer regime, we assume that the switching speed in our architecture is limited primarily by the LC time constant instead of the RC time constant, due to R decreasing as tile size and trace length shrinks:

$$t_{\text{rise}} \sim \sqrt{LC} \quad \text{rather than} \quad t_{\text{rise}} \sim RC$$

To estimate performance of tile-scale interconnects under geometric scaling, we approximate the key electrical parameters using simplified analytical models. These derivations assume uniform tile side length s , a coverage factor α representing the fraction of surface area occupied by traces or pads, and a gap d between tiles.

4.3.1 Capacitance

We approximate the parasitic capacitance between two vertically adjacent tiles (also known as broadside-coupling [58]) using the parallel plate capacitor model, modified by a coverage factor:

$$C = \frac{\epsilon_0 \epsilon_r A}{d} \quad \text{with} \quad A = \alpha^2 s^2$$

$$\boxed{C = \frac{\epsilon_0 \epsilon_r \alpha^2 s^2}{d}}$$

where:

- C : Capacitance between tiles [F]
- s : Tile side length [m]
- α : Surface coverage factor (fraction of tile face covered by conductive traces or pads)




Figure 4.14: modeling parasitic capacitance between two vertically neighboring tiles (in green), which may have different, partially overlapping routing patterns

- d : Gap between adjacent tiles [m]
- A : Effective overlapping conductive area between tiles, $A = \alpha^2 s^2$ [m²]
- ε_0 : Vacuum permittivity ($\approx 8.85 \times 10^{-12}$ F/m)
- ε_r : Relative permittivity of dielectric between tiles (dimensionless)

The surface coverage factor α represents the fraction of each tile face covered by conductive material, including traces and pads. α^2 is used in the capacitance calculation to include the statistical expectation of overlap between two independent conductive patterns on vertically neighboring tile surfaces.

If each surface has αs^2 of metallized area, then the expected overlapping area between two randomly oriented or anisotropically routed layers (e.g. orthogonal traces) is:

$$A_{\text{overlap}} = \alpha^2 s^2$$

This arises from the probability that a point on one surface is metallized (α) and that the corresponding point on the opposing surface is also metallized (α), resulting in $\alpha \cdot \alpha = \alpha^2$ of the total area.

In contrast, if the routing patterns are perfectly aligned (e.g. solid pads on both sides or repeated structures), a single factor α may suffice:

$$A_{\text{overlap}} = \alpha s^2$$

However, realistic 4BI designs may see a variety of routing tiles being used; the uncorrelated or orthogonal layout assumption is more accurate. Thus, the model adopts:

$$C = \frac{\varepsilon_0 \varepsilon_r \alpha^2 s^2}{d}$$

An additional modifier to the original parallel plates capacitance equation requires looking closer at the tile stackup. To realistically map this model to 4BI, the dielectric stackup between plates becomes more complex. As shown in Figure 4.14, the closest valid

geometric arrangement features an air gap and solid substrate (typically FR4 or silicon) between the plates, properly isolating the nets. This arrangement can be modeled as a pair of capacitors in series:

$$\frac{1}{C_{\text{total}}} = \frac{1}{C_1} + \frac{1}{C_2} = \frac{d_1}{\varepsilon_0 \varepsilon_{r1} A} + \frac{d_2}{\varepsilon_0 \varepsilon_{r2} A}$$

where:

- d_1, ε_{r1} : thickness and relative permittivity of the first dielectric (e.g. air)
- d_2, ε_{r2} : thickness and relative permittivity of the second dielectric (e.g. substrate)
- $A = \alpha^2 s^2$: effective overlapping area

This leads to the compact form:

$$C = \frac{\varepsilon_0 \alpha^2 s^2}{\frac{d_1}{\varepsilon_{r1}} + \frac{d_2}{\varepsilon_{r2}}}$$

Practically, this model has implications on channel routing especially when layers start to stack. For example, avoiding overlap by staggering parallel traces on adjacent layers, or minimizing overlap by running traces orthogonal to each other. Additionally, adding shielding through ground pours between layers can help confine capacitive coupling.

4.3.2 Inductance




Figure 4.15: modeling parasitic inductance loop through the smallest possible routing loop between 2x2 tiles

For inductance, we model the loop formed by a signal trace and its return path (via neighboring tiles or vias) as a square current loop of side length s :

$$L \approx \mu_0 \cdot s \cdot \left[\ln \left(\frac{2s}{r} \right) - 0.774 \right] \Rightarrow \boxed{L \sim \mu_0 \cdot s}$$

where:

- L : Inductance of signal-return loop [H]
- μ_0 : Vacuum permeability ($= 4\pi \times 10^{-7}$ H/m)
- s : Tile side length (loop length) [m]
- r : Effective trace width or conductor radius [m]

4.3.3 Resonance Frequency

The natural resonance frequency of the LC circuit formed by the parasitic elements is:

$$f_0 = \frac{1}{2\pi\sqrt{LC}} \sim \frac{1}{2\pi} \cdot \sqrt{\frac{d}{\mu_0\epsilon_0\epsilon_r\alpha^2s^3}} \Rightarrow \boxed{f_0 \propto \frac{1}{s^{3/2}}}$$

where:

- f_0 : Resonance frequency of LC circuit [Hz]
- L : Inductance of signal-return loop [H]
- C : Capacitance between tiles [F]

4.3.4 Capacitive Switching Energy

Assuming a voltage swing V per logic transition, the energy dissipated in charging the capacitance is:

$$\boxed{E_{\text{cap}} = \frac{1}{2}CV^2 = \frac{1}{2} \cdot \frac{\epsilon_0\epsilon_r\alpha^2s^2}{d} \cdot V^2}$$

where:

- E_{cap} : Capacitive energy dissipated per logic transition [J]
- C : Capacitance between tiles [F]
- V : Voltage swing per logic transition [V]

4.3.5 Projections from 10mm down to 10nm

Using the following constants:

- $\epsilon_0 = 8.85 \times 10^{-12}$ F/m — Vacuum permittivity
- $\mu_0 = 4\pi \times 10^{-7}$ H/m — Vacuum permeability
- $\epsilon_{r1} = 1.0$ — Relative permittivity of air
- $\epsilon_{r2} = 3.0$ — Relative permittivity of substrate (e.g. FR4)
- $\alpha = 0.5$ — Surface coverage factor
- $V = 1.0$ V — Voltage swing

And sweeping these equations from $s = 10\text{mm}$ to 10nm , we see the following:
For keeping tile dimensions proportional and shrinking tiles:




Figure 4.16: Parasitic Parameters vs Tile Size ($d_1 = d_2 = 0.33 \cdot s$, Two-Dielectric Model): Capacitance vs Tile Size (a), Inductance vs Tile Size (b), LC Product vs Tile Size (c), Capacitive Energy vs Tile Size (d)

- Capacitance (C) scales linearly with s :

$$C \propto \frac{s^2}{d} \propto \frac{s^2}{s} = s$$

Shrinking tiles reduces capacitance proportionally.

- Inductance (L) also scales linearly with tile size:

$$L \propto s$$

This follows from loop inductance scaling with geometric size.

- LC Product (LC) scales quadratically:

$$LC \propto s \cdot s = s^2 \quad \Rightarrow \quad f_0 \propto \frac{1}{\sqrt{LC}} \propto \frac{1}{s}$$

Resonance frequency increases linearly with decreasing tile size.

- Capacitive Switching Energy (E) scales linearly:

$$E = \frac{1}{2} CV^2 \propto s$$

Resulting in reduced energy per switch at smaller tile sizes.




Figure 4.17: Parasitic Parameters vs Tile Size, Sweep over Coverage ($d_1 = d_2 = 0.33 \text{ } \mu\text{s}$, Two-Dielectric Model): Capacitance vs Tile Size (a), Inductance vs Tile Size (b), LC Product vs Tile Size (c), Capacitive Energy vs Tile Size (d)

These trends support the conclusion that miniaturization improves signal bandwidth and reduces energy dissipation, with no parasitic term that worsens under scale in this model. The dielectric model reflects realistic fabrication constraints and validates the use of LC-dominated timing models for high-speed, small-scale tile architectures.

If we vary surface coverage by sweeping α from 0.1 to 0.9:

We see the following:

- Capacitance (C) increases quadratically with coverage:

$$C \propto \alpha^2$$

This reflects the statistical overlap area between partially metallized tiles. Increasing trace or pad area significantly raises parasitic capacitance.

- Inductance (L) is independent of α :

$$L \propto s$$

Because loop inductance is set by geometric scale and return path, not coverage.

- LC Product (LC) increases with α^2 :

$$LC \propto \alpha^2 s^2 \quad \Rightarrow \quad f_0 \propto \frac{1}{\alpha s}$$

As coverage increases, resonance frequency decreases, setting a tradeoff between signal density and bandwidth.

- Capacitive Switching Energy (E) also increases with α^2 :

$$E = \frac{1}{2}CV^2 \propto \alpha^2$$

Denser routing incurs a power cost, suggesting a need to balance coverage with energy efficiency.

These results demonstrate that although higher α improves connectivity and potentially robustness, it comes at the cost of slower transitions and greater energy consumption. Optimal tile designs must consider α as a critical tuning parameter for performance and power tradeoffs, for example through Joule heating.

Finally, we sweep dielectric layer thickness $d_1 = d_2 = k \cdot s$ from $k = 0.1$ to $k = 0.33$:




Figure 4.18: Parasitic Parameters vs Tile Size, Sweep over $d_1 = d_2 = [0.1-0.33] \cdot s$, (Two-Dielectric Model): Capacitance vs Tile Size (a), Inductance vs Tile Size (b), LC Product vs Tile Size (c), Capacitive Energy vs Tile Size (d)

This reveals that:

- Capacitance (C) decreases as dielectric thickness increases:

$$C \propto \frac{1}{\frac{d_1}{\epsilon_{r1}} + \frac{d_2}{\epsilon_{r2}}} \propto \frac{1}{d}$$

With both layers scaling proportionally to s , increasing k lowers capacitance.

- Inductance (L) remains unaffected:

$$L \propto s$$

Inductance is governed by tile geometry and is independent of dielectric spacing.

- LC Product (LC) decreases:

$$LC \propto C \cdot L \Rightarrow f_0 \propto \frac{1}{\sqrt{LC}} \uparrow$$

Thicker dielectric layers lead to higher resonance frequencies.

- Capacitive Switching Energy (E) also decreases:

$$E = \frac{1}{2}CV^2 \Rightarrow E \propto \frac{1}{d}$$

Lower capacitance leads to reduced energy dissipation per logic transition.

These results confirm that increasing dielectric thickness improves signal bandwidth and reduces energy consumption. However, this must be balanced against mechanical constraints and overall system height. In our 4BI implementations thus far, thickness of substrate significantly affects the complexity of circuit due to height limits of the end effector.

These models form a basis for projecting interconnect performance as tile sizes are scaled down from millimeter, to micrometer, to nanometer regimes, and future work will be done on numerical simulations and experiments to validate these analytical performance projections.




Figure 4.19: Routing considerations: parallel layer 1 (a), parallel layer 2 (b), crossed layer 1 (c), crossed layer 2 (d)

5 4BI Connector (4BIc)

4BI has clearly demonstrated its ability to be automatically assembled by an assembler and still has a lot of headroom to explore. However, joint reliability remains an issue; scaling is pointless if the approach itself is flawed.

While 4Hp was difficult to scale, its H-eon interconnects made for very reliable connections. How can we incorporate these features from 4Hp into 4BI?

The answer is 4BIc, incorporating a leadframe into the tile itself adds compliance at each joint, which evenly distributes forces and creates conductivity at reasonable forces.

While I spent most of my time on VMDs and DICE, I also spent time developing a process flow for fabrication of electrical Connectors in a Day (eCiD).

5.1 Electrical Connectors in a Day (eCiD)

Over the past two decades, additive and subtractive desktop fabrication technologies have steadily improved in reliability and resolution, to the point where the edge of their capabilities are able to prototype electrical connectors, which may require features in the 10s of microns. Sub-\$5000 desktop machines such as the Xtool F1 Ultra, Carvera CNC, and Formlabs Form 4 are able to achieve resolutions in this ballpark, and example interconnects were fabricated to explore this capability space.

Connector Type	Fabrication Method
Header 1xn	Leadframe: Fablight, Carvera CNC
IDC 0.1" 2x2	Leadframe: F1 Ultra
SlimStack 0.8mm	Leadframe: F1 Ultra
DICE 4Cx tile	Leadframe: F1 Ultra
H-interconnect (eye-of-the-needle)	Leadframe: Fablight

5.1.1 Assembly Techniques

Towards making connector assembly techniques more accessible, I developed several proxy approaches to approximate an insert-molded or contact loaded interconnect, shown in Figure 5.2. In some cases, the proxy was exactly the same approach as would be traditionally done, except I used digital fabrication techniques to replace conventional manufacturing methods, such as substituting an injection-molded housing with a resin printed part.

For insert molded proxies, I explored two primary techniques. The clam-shell technique involved splitting housings into two parts (part A and B), where the leadframe was

Connector-in-a-Day


Header 1x4



IDC 0.1" 2x2



SlimStack 0.8mm



DICE 4Cx




Figure 5.1: Connectors fabricated using a combination of desktop fabrication tools

mated to registration features between the two parts. After adding adhesive, I used the two parts to sandwich and seal the leadframe. This approach allowed me to achieve precision alignment without specialized molding equipment. I also investigated the print-pause technique (insert-printing), where I started an FDM print using a 0.2mm nozzle (or smaller), paused the operation to insert a leadframe, and then continued the print. This method effectively embedded the leadframe within the printed structure, creating a robust mechanical and electrical assembly.

For situations where traditional contact loading was preferred, I developed a loaded terminal proxy approach. In this process, I printed a housing instead of injection molding it, but then loaded contacts normally as would be done in conventional connector assembly. This hybrid approach maintained compatibility with existing connector designs while leveraging the accessibility of digital fabrication for the housing components.

Assembly techniques




Figure 5.2: Assembly techniques

5.1.2 Housings

Connector housings typically rely on micromolding and operate very close to the resolution limits of desktop mSLA printers, presenting unique challenges when supporting small prints. Conventional SLA support structures fail to provide even surface support across delicate features, often resulting in inconsistent quality. While line supports offer improved structural integrity, they frequently risk damaging fragile features during removal—an issue that persisted despite implementing tapered interfaces to minimize contact area.

Surprisingly, I observed no significant performance difference between horizontal versus vertical line orientations. The printing process requires meticulous tuning for small features, as achieving good surface finishes proves exceptionally difficult, and separating delicate parts from the print bed without damage presents a persistent challenge.

To address these limitations, I developed suspended tab supports that distribute forces evenly across small features (including walls as thin as 0.1mm), enabling clean break-off without compromising part integrity. This approach significantly simplifies print bed removal, as conventional supports now attach to these sacrificial tabs rather than directly to the functional part, providing a reliable pathway to fabricate the fine connector geometries essential to fabricating VMD geometries.

Additionally, I discovered that adding backside channels to alleviate cupping (Figure 5.4) from internal contour features makes a significant improvement to dimensional accuracy of prints.

5.1.3 Leadframes

Leadframes are usually made out of red metals (copper-based) to maximize electrical conductivity. C110 was a good starting point, but for fabricating compliant flexures, I preferred Phosphor Bronze (C510) or Brass.

In my approach, I used fiber lasers to cut leadframe positives out of stock sheet material. I worked with both the machine shop 3kW Fablight and the desktop Xtool 20W F1 Ultra, demonstrating the process at different scales. For relatively thicker metals, like 0.6mm, I preferred using the Fablight. For thinner metals, around 0.127mm, I preferred using the F1 Ultra.


My leadframe designs incorporated several standard features. I included pilot holes to register the carrier to the process. In high-volume reel-to-reel processes, these pilot holes would continuously register to the conveyer motion system, guiding materials through multiple processes progressively. This enables techniques such as progressive stamping.

I implemented tiebars to mechanically hold and support design features in place during processing. Because they needed to mechanically support design features, I found the orientation of the leadframe with respect to gravity and direction of travel was important for retaining geometry shape; lack of sufficient support easily led to deformed features. This was especially important in a high-volume reel-to-reel context; while certain designs could retain their shape at low speeds, high speeds may introduce additional forces.

After singulation (separation of tiebars), contacts were typically the end result of my leadframe fabrication process.

Printing small connectors - challenges


- Conventional SLA supports don't evenly support surfaces
- Line supports improve support, but risk breaking fragile features
- Difficult to get good surface finishes
- Tuning for small features
- Difficulties removing from print bed



(a)

Printing small connectors - solutions

- Suspended tab supports evenly support small features (~0.1mm walls)
- Easy break off without damaging part
- Easy to remove from print bed (conventional supports attach to tabs)



(b)

Figure 5.3: Printing supports for small connector housings




Figure 5.4: Backside channels to alleviate cupping

For tooling, I used forming to deform leadframe contacts out of plane. This helped me accomplish many objectives, but the primary one was making sure the contact interfered with its mating target for electrical conductivity. I also explored coining, a variation on forming where a small precise feature of the metal was deformed out of plane, also useful for creating interference with a mating target. Finally, I used singulation to shear away the metal, typically for removing tiebars and isolating the contacts in a leadframe.

IDC leadframe - flow




Figure 5.5: Leadframe flow for the IDC

5.1.4 Process Registration

Process registration is necessary for multi-process operations where the workpiece needs to be moved from process to process. Registration is typically carried out with dowel pins, which are precision ground steel rods with tapered ends. Mating features (pilot holes) are included on leadframes.

I chose to use 2mm undersized dowel pins, which were sufficient for most designs and save on design space. They can be used in conjunction with different sized dowel pins to differentiate different purposes.

Complying with kinematic principles, 3 points of contact are sufficient for constraining a plane. 4 will overconstrain and lead to deformation of the plane.

IDC - form tool




Figure 5.6: Form tool for the IDC

IDC - form tool




Figure 5.7: Fabricated tools for the IDC

5.2 4BIc Geometry




Figure 5.8: 4BIc dimensions: bare 4BIc0 tile dimensions (a), 4BIc tile with contacts (b), side view of the contact geometry (c)

4BIc is the latest and greatest 4Bx geometry, designed for batch production to scale towards the 1000s of tiles required for bigger circuits, like the RISC-V processor on the order of 10,000s of tiles. It improves on 4BI by adding a leadframe, creating “leaf spring” compliant geometry for each contact. This approach is inspired by the leadframe principles present in connectorized electronics packaging technologies such as the ASEP technology [59]; [60]; [61].

5.2.1 4BIc Process Steps Overview




Figure 5.9: 4BIc process steps: (a) PCB 5x5 v-cut panel of sot323 tiles, (b) Fiber laser-cut leadframe, (c) Eject trash for solder rivets later, (d) Form 1: 90 degree bend, (e) Eject leadframe, (f) Formed leadframe 1, (g) Flip over, (h) Register against PCB for reflow step, (i) Form 2: 60 degree bend, create compliant contacts, (j) Formed leadframe 2, (k) Laser Singulation: remove leadframe tie-bars, (l) Tool Singulation: shear PCB v-cuts





Figure 5.10: 4BIc assembled into a circuit

5.2.2 Panelized PCBs


Over the course of geometry iteration, having enough tiles to test assembly and build circuits with became a bigger bottleneck than design iterations themselves. For 4BIc, it

made sense to order many tiles, since it had been relatively stable as a design and we've made good progress on assembly with it.


Each panel contains $5 \times 5 = 25$ tiles, and were ordered as v-cut panelized panels. This was deliberately chosen over mouse-bites due to ease of singulation by tool. In an attempt to standardize certain elements of the design, such as silkscreen patterns for visibility, both sides of the PCBs looked nearly identical. Laser marking was used to differentiate the backside from the front, conveniently a 2 second job.




(a) dowels



(b) dowels2



(c) f1-ultra-fixtured



(d) soldermask-array

Figure 5.11: Laser-marking soldermask for differentiating sides

5.2.3 Leadframe Fabrication

4BIc uses 0.127mm thick phosphor bronze (C510) stock for its leadframe, which takes ~4-5 minutes to cut on the F1 Ultra (Figure 5.12).

There were challenges maintaining quality of cut for leadframes, and especially when considering fabricating more at small batch production. Often, the leadframe or trash would weld itself to the steel backing, which involved scraping to remove and clean. I found polyimide backing tape was the solution to this persistent issue, significantly improving both yield and processing time.

Tape fixturing orientation proved important when dealing with significant warping, which could cause stock to lift away from the focal plane during laser processing. Through experimentation, I determined that more tape generally produced better results, particularly when strategically placed to counter material warping tendencies.

Conveniently, I discovered that 0.127mm phosphor bronze was thin enough for efficiently cutting with scissors and paper guillotines, instead of power shears. This unexpectedly simplified certain processing steps, especially for quick modifications or when separating smaller pieces from larger stock sheets.




Figure 5.12: 0.127mm C510 leadframe fabrication with the F1 Ultra

5.2.4 Ejecting Trash

After cutting, if there are closed contour features on the inside of the pattern (such as holes), trash can get trapped. This is especially problematic if the trash is small, such as holes for solder rivets. In 4BIc, the leadframe is designed to be “married” to the PCB using a solder rivet approach, which is why ejection needs to be repeatable and comprehensive. Form 4 Rigid 10k was used to print most of the tooling, including the ejection tool initially used to remove trash.


Later, this problem was solved by reducing the number of closed countour features and merging them with larger contours, since larger trash is easier to remove.

5.2.5 Form 1 Tooling for Leadframe


A lot of iteration was given to the first form (1) process. The features happened to be small enough, and the tool lacked the proper registration lead-in features that forming caused breakage in the tool. The failure mode suggested misalignment (bias to one side); the first tool lost only teeth in columns to one side.

I implemented several key improvements to address these issues. I ended up adding drafted lead-in features that engaged first to figure out alignment before the tool itself engaged, creating a self-centering effect that significantly reduced misalignment failures. Additionally, I merged the tool features together to create bigger, more robust structures, which improved durability of the tool during repeated forming operations. I also observed that the leadframe itself was twisting more than expected, primarily due to thin tie-bars, so I redesigned these elements to be thicker, effectively reducing unwanted deformation during the forming process. These modifications collectively transformed an initially unreliable forming operation into a repeatable manufacturing step essential for producing consistent leadframe geometries.


A significant feature for correctly forming parts and avoiding misalignment is adding lead-in; this feature should engage with the other half of the tool before fragile features




(a) eject-tool-post



(b) holes2




(c) tool-installed



(d) setup




(e) arbor




(f) good-and-bad


Figure 5.13: Tool assembly process: the print is post-cured in the Form Cure to maximum strength properties (a), parts, particularly the mating clamp, are checked for clear channels (b), tools are installed into their receptacles (c), the top half is then loaded onto the linear guides (d), the tool is loaded into the arbor press (e), and then the press is actuated and the leadframe has been modified, whether ejected/formed/singulated




(a) broken-teeth-4




(b) clamp




(c) forming-leadframe




(d) leadframe-dowels




(e) leadframe-formed



(f) top-tool




(g) tough-1500




(h) twisted-tiebars

Figure 5.14: The Form tooling process

in the tools engage. In a way, this lead-in acts like a crossing guard, making sure the two halves are aligned before crossing the street.



(a) misalignment



(b) registration-lead-in

5.2.6 Formed Leadframe

As mentioned earlier, inside contour features were merged; the holes used for solder rivets were merged with the large center trash. This design modification produced several beneficial effects. The leadframe became less affected by heat, potentially causing less warping during thermal processing steps. Additionally, the trash became much easier to extract, and could be done by hand rather than requiring a specialized tool. These seemingly minor geometric adjustments significantly streamlined the manufacturing process while improving overall part quality and consistency.


5.2.7 Form 2 Tooling for Leadframe and Pcb

Unfortunately, form 1 doesn't appear to form the contacts far enough. The form 2 operation has trouble engaging the leadframe and forming it in the right direction. Also, because form 2 has slanted features, the layer lines add roughness to the operation, which can cause binding of the fragile contacts. One solution would be to apply this form one direction at a time, requiring four applications. Another solution, potentially merging the operations from form 1 and form 2 into a single step could also solve this problem; this iteration will have to be future work.


5.2.8 Tool Singulation of Pcb Panel

Although the F1 Ultra can be used to singulate FR4, it made more sense to singulate by tool. V-cut panelization typically leaves too many hairs if singulated by hand, and can be arduous (a previous timing using snippers took nearly 20min!). However, like the previous steps, a singulation tool can streamline the process, with clean and repeatable results; this tool takes up to 60 seconds to operate, if there is a jam.


Initially, the tool was made to punch all tiles out simultaneously (Figure 5.18a), upon reflection I changed the tool to progressively punch the tiles out to reduce forces (Fig-



(a) formed




(b) form1a




(c) form1b

Figure 5.16: Formed leadframe, with revised geometry solving earlier problems



(a) form2-tool




(b) alignment




(c) misalignment

Figure 5.17: The second form operation, which uses slanted tooling features to form the final leaf spring contacts



(a)




(b)




(c)

Figure 5.18: Tile singulation tool


ure 5.18b). Like the other steps, there is a top (tool) half, and a bottom (clamp) half. These parts are also printed from Rigid 10k (Figure 5.18c).




(a) sing-pattern



(b) increase-opening



(c) edge-quality



(d) discarded-frame

Figure 5.19: Singulation troubleshooting: The first singulation left a pattern that correlates with the first wave of punches (a), the rest were stuck in the clamp (b). Increasing the opening of each of the windows resolved the issue. After ~20 singulations, edge quality seems to be fairly consistent (e).

5.2.9 Advancing VMDs

The 4B1c geometry represents a significant advancement in the VMD ecosystem, addressing key limitations identified in previous geometries while establishing a scalable manufacturing pathway for complex circuit assembly. By integrating compliant lead-frame elements directly into the tile structure, 4B1c successfully balances mechanical reliability with electrical performance – the critical challenge that constrained earlier VMD implementations.

Throughout my evaluation of joint reliability and circuit performance, I identified several persistent issues with the 4BI approach, particularly regarding intermittent electrical contacts and the high normal force window required for stable operation. Traditional pad-to-pad connections proved vulnerable to mechanical shock, vibration, and inconsistent pressure distribution, which became increasingly problematic when scaling to larger circuit assemblies. 4BIc directly addresses these challenges by incorporating leaf-spring contacts that provide consistent, reliable electrical connections at significantly lower activation forces while accommodating minor misalignments.

The manufacturing process I developed for 4BIc, though still evolving, demonstrates a viable path toward high-volume production. My stepwise approach – from panelized PCB fabrication through leadframe cutting, forming, and singulation – is deliberately designed to leverage accessible digital fabrication technologies while maintaining compatibility with established electronics manufacturing techniques. This hybrid strategy allows 4BIc to serve as a transitional technology that can be implemented immediately with today’s tools while demonstrating the architectural principles that will define future VMD systems.

Looking forward, several improvements to the 4BIc implementation could yield significant benefits. Integrating form 1 and form 2 operations into a single step would streamline manufacturing, while further refinement of the leadframe design could improve compliance characteristics. Experimenting with alternative plating finishes for the contact surfaces may reduce contact resistance even further, potentially enabling higher-frequency applications. Most importantly, now that the assembly system and manufacturing process have been established, extensive reliability testing of completed 4BIc circuits will provide crucial data for optimizing future designs.

By addressing the fundamental reliability challenges identified in earlier VMD geometries, 4BIc paves the way for scaling beyond simple proof-of-concept circuits to the larger, more complex assemblies required for meaningful computational applications. The completed RISC-V processor envisioned in this work, requiring tens of thousands of tiles, now stands as an achievable milestone on the path toward truly volumetric electronics packaging.

6 Conclusion

“There’s always room at the middle.”

— Suraj Bramhavar,
ARIA Meeting, May 2025

This thesis introduces the volume mount device (VMD) as an alternative to surface-mount device (SMD) standards, reimagining electronic packaging for true three-dimensional assembly rather than conventional two-dimensional integration.

The VMD framework embeds electrical function and mechanical structure within modular elements that form self-constraining 3D lattices. While currently incorporating SMD components on tile PCBs, this approach establishes a pathway toward eventually replacing SMDs at the IC packaging level.

My hybrid assembly system advances beyond prior work by placing hundreds of integrated elements at ~1000 CPH, compared to previous systems limited to tens of components at slower rates.

My evaluation of geometric configurations, performance overheads, and self-aligning connector interfaces establishes VMD as a viable approach for rapid electronics prototyping. Beyond simply replacing breadboards and prototype PCBs, this work lays the foundation for a scalable, reversible, and reconfigurable packaging framework inspired by natural ecosystems’ circularity. The VMD approach offers a compelling path toward more sustainable electronics systems that can be assembled, disassembled, and reassembled as needed—reducing waste while improving resilience against supply chain disruptions.

As electronics continue to evolve, the volumetric integration method demonstrated in this thesis provides a promising direction for overcoming the limitations of planar manufacturing while leveraging existing supply chains in new, more flexible ways.

6.1 Future Work

Having demonstrated an automated system capable of assembling 100s of integrated electronic digital materials, there are a few next steps that can be undertaken.

6.1.1 Finishing 4Blc

4Blc is a few steps away from being finished, but was not finished in time for this thesis. The leadframe tooling is almost there, and a new approach may consolidate multiple form steps into one, simplifying assembly and increasing reliability. Considering the results from the full-adder evaluation, improving joint reliability is a top priority. Multiple solutions have been suggested towards evenly distributing forces between contacts, such

as adding a compliant conformal layer to the cap itself, but adding compliant contacts to each pad theoretically remains the best solution that not only guarantees even distribution of force between the cap and the assembly, but predictable forces throughout the entire volume. This solution also remains the most scalable, with ample headroom for reaching higher orders of magnitude.

Joint evaluation tests will also have to be conducted on 4B1c, such as Normal Force Window and cycling tests. In addition, to validate our analytical models, numerical and experimental work will need to be done to characterize capacitive and inductive parasitics, for evaluating switching frequency as feature sizes shrink.

Breaking out connections using solder joints and pyralux was a source of error; it partially relieved strain from wires directly soldered to tiles, but could be improved (experimenting with thinner wires with less strain, applying principles from 4B1c and connector fabrication to create compliant test probes, etc.).

6.1.2 Assembly Improvements

Several improvements could immediately enhance the assembly process. For example, close-looped feeding would enable the system to effectively recycle tiles without manual intervention, reducing operator workload for large-scale assemblies. Scaling to larger assemblies will necessitate this feature; disassembly will become more costly than assembly itself.

Kinematic fixtures would reduce the need for recalibration moving feeders and templates on and off the assembler bed, saving a significant amount of time.

Additional computer vision could be used to add additional error checks, reducing the need for manual intervention and improving overall system reliability. For gathering longer term reliability data, it may start making sense adding additional cameras monitoring other aspects of the assembly process.

6.1.3 Geometric Exploration for Enhanced Reliability

While square geometries provide compatibility with existing supply chains and sufficient I/O capabilities, future work should explore alternative geometric configurations that enhance mechanical reliability. Although triangular elements offer superior kinematic stability with their three points of contact, they present challenges for electronic routing due to limited I/O capacity.

A promising research direction would be hybrid geometries that combine triangular mechanical contact points with square or hexagonal electrical interfaces. This approach could yield self-aligning structures with guaranteed contact between mechanically independent elements while maintaining adequate routing capabilities for power, ground, data, and clock signals.

Additionally, exploring true 3D elements (beyond stacked 2D elements) could unlock unique advantages in tessellation and self-orientation behavior, though this would require addressing the additional fabrication complexity outside the current manufacturing mainstream.

6.1.4 4Cx Family - 4C, 4Ci, 4Ce, 4Cs

In addition to 4Bx work, efforts were made towards exploring more self-aligning geometries in the 4Cx (connector¹) family, which are also capable of handling higher pin count and higher power necessary for application goals such as variable-pin chiplets and power electronics, shown in Figure 6.1.

Initial work on 4C relied on chamfered edges mating flush with adjacent tiles. However, the design of these edges only constrained the tiles in height, and not laterally in xy. 4Ce (edge) introduced “kinematic” french-cleat like designs inspired by Jake Read’s “Kleat Toolchanger” mechanism from his Clank machines project [62]. My implementation in geometry enabled constraint in both height and xy, but created failure modes where the edge could act as a pivot point, destabilizing the structure when subject to external forces. Finally, 4Cs (surface) introduced aligning pillars inspired by 4BI, but directly on the surface of the tiles.

In these approaches, devices are integrated directly into the tile rather than sit on the surface, simplifying the pick and place problem for arbitrary devices of variable pin count.

These approaches represent experiments towards shifting the geometry feedstock itself away from conventional 2d approaches, such as PCBs, towards more 3d approaches, closer to Molded Interconnect Devices (MIDs) [63] and ASEP [59]; [60]; [61].

6.1.5 uVMDs and nVMDs

While this thesis focuses primarily on scaling element count using millimeter-scale components to enable sophisticated circuits, a parallel and complementary trajectory involves scaling down feature sizes from millimeter to micrometer regimes. This miniaturization approach not only allows more elements to fit within a given volume but also significantly reduces performance overhead and parasitic effects.

My colleague Teddy Hsieh from the Niroui Group is currently fabricating nano and micro devices using discrete feedstock such as 55nm gold nanoparticles, building upon Spencer Zhu’s foundational work in this area. Their research group has established a strong track record in pushing the boundaries of nanoscale engineering [64] [65] [66]. Our collaborative roadmap aims to bridge the hierarchical assembly and automation techniques developed in this thesis with Teddy’s device fabrication expertise, creating a unified approach that spans from nano to macro scales.

Accessing these smaller regimes requires specialized fabrication tools. The UpNano One 2-photon printer represents a promising approach for rapid prototyping of micro-scale elements. Operating similarly to a desktop SLA printer but with significantly higher resolution, this system enables direct fabrication of structures in the micrometer or even nanometer range.

This miniaturization approach shifts the VMD paradigm closer to electronics packaging level 1 (device-in-IC) rather than level 2 (SMD-on-PCB). Beyond the obvious performance advantages from reduced parasitic effects, this scale reduction offers practical

¹the “C” in 4Cx could be interpreted as the shape of a compliant contact, in keeping with the naming scheme




Figure 6.1: 4Cx family of geometries







Figure 6.2: ~55nm gold nano cubes undergoing contact printing for precise, scalable, and pristine particle patterning, [64]




(a) overview




(b) 4mm elements



(c) 400um elements



(d) 4mm elements, up close



(e) 400um elements, up close

Figure 6.3: 4Cx VMD prototypes (with pockets for integrating components) printed on the UpNano One 2-photon printer

assembly benefits. Smaller elements enable more compact circuits, making parallel processing with multiple assembly machines more feasible by reducing the physical workspace requirements for complex structures.

6.1.6 Mini-Assembly Farms for Smaller Elements

Improving serial performance vs building parallel machines was discussed in Chapter 3. However, another interesting angle borrows from the abstraction levels introduced earlier for digital materials: what is the minimal form of a machine? My lab has been focused on this question for awhile, deriving inspiration from biology and working towards relative machines [47]; [67].

While the assemblers assembling assemblers arc is important this research space, it is fairly ambitious in the near term, and would be difficult to immediately apply to my geometry. A close analogy might be a minimal conventional machine that could be rapidly fabricated, assembled, and deployed. By reducing the complexity and footprint of individual assembly units, this approach could potentially enable higher density deployment of assembly capacity within a given workspace.

Critical to the success of assembly farms will be solving the challenges of machine-to-machine transport and registration. Effective systems must establish reliable protocols for transferring partially completed assemblies between machines while maintaining precise alignment. This would require advanced localization methods, potentially utilizing computer vision, fiducial markers, or mechanical registration features to ensure seamless coordination between independent assembly units working on different aspects or regions of a single construction.


The incremental improvements discussed earlier in the thesis combined with this more ambitious assembly farm concept provide a roadmap for scaling VMD assembly capacity by orders of magnitude beyond what has been demonstrated in this thesis.

6.1.7 Super-DICE 2.0


Super-DICE 2.0 picks up on Zach Fredin and Camron Blackburn’s work on Super-DICE [68] [42], this time working with Camron Blackburn and Alex Wynn; example devices from their work are shown in Figure 6.4.

Super conducting electronics (SCE) have unique packaging challenges. For one, Josephson Junctions (JJs), a fundamental building block for SCE, are fabricated at much coarser resolutions than conventional lithographic transistors, which makes it difficult to integrate sufficient complexity into a single die to enable compute applications (10k JJs vs millions, billions transistors). Modular approaches are therefore of great interest towards moving past single device performance.

Using a modified 4BI approach, the plan is to assemble these devices. Devices will either be packaged onto PCB substrates as typical, or the bare die themselves can act as the tiles. This application is particularly interesting for 4BI, as it represents a major step towards the integrated device abstraction level, with higher i/o per interface (Figure 6.4b).



(a)



(b)

Figure 6.4: Super-DICE 2.0: AQFP buffer device designed by Camron Blackburn et al [69] (a), Concept layout for Super-DICE 2.0, Camron Blackburn's sketch (b)

6.2 Impact

The Volume Mount Device framework represents a fundamental reimagining of electronic systems integration, challenging the two-dimensional paradigm that has dominated the industry for decades. VMDs enable more resilient electronics that can withstand supply chain disruptions while opening design possibilities previously constrained by planar thinking. As the electronics industry pushes against the physical limits of traditional approaches, the volumetric integration methods demonstrated here offer a valuable complementary trajectory. Just as the transition to surface-mount technology revolutionized electronics manufacturing in the late 20th century, volume-mount devices may define the next generation of electronics integration.

Bibliography

1. Bohr M (2007) A 30 Year Retrospective on Dennard's MOSFET Scaling Paper. IEEE Solid-State Circuits Newsl 12(1):11–13. <https://doi.org/10.1109/N-SSC.2007.4785534>
2. Leiserson CE, Thompson NC, Emer JS, et al (2020) There's plenty of room at the Top: What will drive computer performance after Moore's law? Science 368(6495):eaam9744. <https://doi.org/10.1126/science.aam9744>
3. The Death of Moore's Law: What it means and what might fill the gap going forward | CSAIL Alliances. <https://cap.csail.mit.edu/death-moores-law-what-it-means-and-what-might-fill-gap-going-forward>. Accessed 13 May 2025
4. Is Moore's law dead? <https://www.imec-int.com/en/semiconductor-education-and-workforce-development/microchips/moores-law/moores-law-dead>. Accessed 13 May 2025
5. Ghani T, Ranade P (2024) [The Incredible Shrinking Transistor - Shattering Perceived Barriers and Forging Ahead](#). In: 2024 IEEE International Electron Devices Meeting (IEDM). pp 1–4
6. published AS (2025) Intel has championed High-NA EUV chipmaking tools, but costs and other limitations could delay industry-wide adoption: Report. <https://www.tomshardware.com/tech-industry/intel-has-championed-high-na-euv-chipmaking-tools-but-costs-and-other-limitations-could-delay-industry-wide-adoption-report>. Accessed 13 May 2025
7. Xie C, Zhang Y, Zhang K, Guo LZ-R (2023) Introduction of Gate-All-Around FET (GAAFET). ACE 28(1):164–175. <https://doi.org/10.54254/2755-2721/28/20230354>
8. Intel Is All-In on Backside Power Delivery - IEEE Spectrum. <https://spectrum.ieee.org/backside-power-delivery>. Accessed 13 May 2025
9. This Machine Could Keep Moore's Law on Track - IEEE Spectrum. <https://spectrum.ieee.org/high-na-euv>. Accessed 20 May 2025

10. (2024) 5 things you should know about High NA in EUV. <https://www.asml.com/en/news/stories/2024/5-things-high-na-euv>. Accessed 20 May 2025
11. (2021) What is EUV lithography? <https://research.ibm.com/blog/what-is-euv-lithography>. Accessed 20 May 2025
12. What are Chiplets and Why are Chiplets Important? Chiplets Overview |Cadence. https://www.cadence.com/en_US/home/explore/chiplets.html. Accessed 13 May 2025
13. Heterogeneous Integration Roadmap - IEEE Electronics Packaging Society. <https://eps.ieee.org/technology/heterogeneous-integration-roadmap.html>. Accessed 13 May 2025
14. (2021) Paul Eisler - Inventor of the printed circuit board. <https://ats.net/en/paul-eisler-inventor-of-printed-circuit-board/>. Accessed 16 May 2025
15. Why Is There a Chip Shortage? Covid-19, Surging Demand Cause Semiconductor Shortfall. <https://www.bloomberg.com/graphics/2021-semiconductors-chips-shortage/>. Accessed 12 May 2025
16. Impact of the New U.S. Tariff Policy on Customs Clearance - News - PCBway. https://www.pcbway.com/blog/News/Impact_of_the_New_U_S_Tariff_Policy_on_Customs_Clearance_51dff4fd.html. Accessed 12 May 2025
17. The world's first 0.1 μ F in 008004 inch size capacitor | Product & Event News. <https://www.murata.com/en-us/news/capacitor/ceramiccapacitor/2019/1205>. Accessed 16 May 2025
18. Murata Unveils the World's Smallest Multilayer Ceramic Capacitor with the First 006003-inch Size (0.16mm $\times$ 0.08mm) Device | Product & Event News | Murata Manufacturing Co., Ltd. <https://www.murata.com/news/capacitor/ceramiccapacitor/2024/0919>. Accessed 16 May 2025
19. Froese M Samsung launches advanced 2.5D chip-packaging technology, I-Cube4. <https://www.engineersgarage.com/samsung-launches-advanced-2-5d-chip-packaging-technology-i-cube4/>. Accessed 14 May 2025
20. Leibson S (2011) Generation-jumping 2.5D Xilinx Virtex-7 2000T FPGA delivers 1,954,560 logic cells, consumes only 20W. <https://web.archive.org/web/20250513210135/https://www.design-reuse.com/industryexpertblogs/27614/2-5d-xilinx-virtex-7-2000t-fpga.html>. Accessed 13 May 2025

21. Jikang L (2025) The application of organic materials used in IC advanced packaging:A review. *Memories - Materials, Devices, Circuits and Systems* 9:100124. <https://doi.org/10.1016/j.memori.2025.100124>
22. Priyadarshi K (2023) What is 2D, 2.5D & 3D Packaging of Integrated Chips? <https://techovedas.com/what-is-2d-2-5d-3d-packaging-of-integrated-chips/>. Accessed 14 May 2025
23. published PA (2023) AMD Shares New Second-Gen 3D V-Cache Chiplet Details, up to 2.5 TB/s. <https://www.tomshardware.com/news/amd-shares-new-second-gen-3d-v-cache-chiplet-details-up-to-25-tbs>. Accessed 20 May 2025
24. Cutress DI AMD Demonstrates Stacked 3D V-Cache Technology: 192 MB at 2 TB/sec. <https://www.anandtech.com/show/16725/amd-demonstrates-stacked-vcache-technology-2-tbsec-for-15-gaming>. Accessed 20 May 2025
25. published PA (2023) AMD unveils Instinct MI300X GPU and MI300A APU, claims up to 1.6X lead over Nvidia's competing GPUs. <https://www.tomshardware.com/pc-components/cpus/amd-unveils-instinct-mi300x-gpu-and-mi300a-apu-claims-up-to-16x-lead-over-nvidias-competing-gpus>. Accessed 20 May 2025
26. Kennedy P (2024) Broadcom 3.5D XDSiP with Face-To-Face 3.5D For 2026 XPU's and Beyond. <https://www.servethehome.com/broadcom-3-5d-xdsip-with-face-to-face-3-5d-for-2026-xpus-and-beyond/>. Accessed 20 May 2025
27. Specifications. <https://www.uciexpress.org/specifications>. Accessed 20 May 2025
28. Ahmad M (2023) A sneak peek at chiplet standards. <https://www.edn.com/a-sneak-peek-at-chiplet-standards/>. Accessed 20 May 2025
29. Synopsys to Acquire Ansys, Creating a Leader in Silicon to Systems Design Solutions. <https://www.ansys.com/news-center/press-releases/1-16-24-synopsys-acquires-ansys>. Accessed 20 May 2025
30. 3DIC Compiler: Platform for Multi-Die Designs | Synopsys. <https://www.synopsys.com/implementation-and-signoff/3dic-design.html>. Accessed 20 May 2025
31. Cadence 3D-IC Design | Proven Design Flows for Multi-Chiplet Design and Advanced IC Packaging. https://www.cadence.com/en_US/home/solutions/3dic-design-solutions.html. Accessed 20 May 2025

32. published AS (2021) Motherboard and GPU Pricing Under Pressure Due to Copper Foil Shortage. <https://www.tomshardware.com/news/motherboard-and-gpus-under-pressure-due-to-copper-foil-pricing>. Accessed 12 May 2025
33. Frequently Asked Questions: CHIPS Act of 2022 Provisions and Implementation. <https://www.congress.gov/crs-product/R47523>. Accessed 13 May 2025
34. David E (2024) Who benefitted from the CHIPS Act? <https://www.theverge.com/24166234/chips-act-funding-semiconductor-companies>. Accessed 24 Nov 2024
35. Salvi A, Vardya A, Whiteside S (2024) CHIPS Act doesn't go far enough to protect America's supply of microelectronics. Here's how to fix it. <https://chicago.suntimes.com/other-views/2024/07/02/chips-act-printed-circuit-board-pcb-technology-electronics-aidan-salvi-anaya-cardya-shane-whiteside>. Accessed 12 May 2025
36. (2024) AI's carbon footprint is bigger than you think | MIT Technology Review. <https://web.archive.org/web/20240705073854/https://www.technologyreview.com/2023/12/05/1084417/ais-carbon-footprint-is-bigger-than-you-think/>. Accessed 20 May 2025
37. Gelles D (2024) A.I.'s Insatiable Appetite for Energy. The New York Times: Climate
38. (2025) Explained: Generative AI's environmental impact. <https://news.mit.edu/2025/explained-generative-ai-environmental-impact-0117>. Accessed 20 May 2025
39. Wire Wrap Is alive And Well! https://www.nutsvolts.com/magazine/article/wire_wrap_is_alive_and_well. Accessed 20 May 2025
40. Wire-Wrap_Ad_Bell_Labs_Oct53.jpg (JPEG Image, 3109 × 4329 pixels) — Scaled (15%). https://bitsavers.org/pdf/att/Bell_Labs/Wire-Wrap_Ad_Bell_Labs_Oct53.jpg. Accessed 20 May 2025
41. 1953 U.S. NAVY ELECTRONIC MANUFACTURING DEVELOPMENT PROGRAM PROJECT TINKERTOY 47174
42. Fredin Z (2021) *Assembling Integrated Electronics*. Thesis, Massachusetts Institute of Technology
43. Langford WK (2014) *Electronic digital materials*. Thesis, Massachusetts Institute of Technology

44. Ward J(JonathanD (2010) [Additive assembly of digital materials](#). Thesis, Massachusetts Institute of Technology
45. MacCurdy R, McNicoll A, Lipson H (2014) Bitblox: Printable digital materials for electromechanical machines. *The International Journal of Robotics Research* 33(10):1342–1360. <https://doi.org/10.1177/0278364914532149>
46. Hiller J, Lipson H (2009) Design and analysis of digital materials for physical 3D voxel printing. *Rapid Prototyping Journal* 15(2):137–149. <https://doi.org/10.1108/13552540910943441>
47. Langford WK (2019) Discrete Robotic Construction
48. Langford W, Ghassaei A, Gershenfeld N (2016) [Automated Assembly of Electronic Digital Materials](#). In: Volume 2: Materials; Biomanufacturing; Properties, Applications and Systems; Sustainable Manufacturing. American Society of Mechanical Engineers, Blacksburg, Virginia, USA, p V002T01A013
49. Brearley D (2016) Great Connector Inventions: Compliant Press-Fit Pins. <https://connectorsupplier.com/great-connector-inventions-compliant-press-fit-pins/>. Accessed 17 May 2025
50. Feeding Technology. <https://www.depragusa.com/Feeding-Technology.html>. Accessed 14 May 2025
51. LumenPnP - Accurate and Reliable Pick and Place Machine – Opulo. <https://www.opulo.io/products/lumenpnp>. Accessed 16 May 2025
52. NeoDen 3V-A Automatic SMD Pick and Place Machine Manufacturers and Suppliers China - Wholesale Products - Neoden Technology. <https://www.neodensmt.com/pick-and-place-machine/neoden-3v-a-automatic-smd-pick-and-place.html>. Accessed 16 May 2025
53. RS-1R: Fast Smart Modular Mounter. <https://jukiamericas.com/pages/rs-1r>. Accessed 16 May 2025
54. NXT III Placing Machines | Fuji America Corporation. <https://www.fujiamerica.com/NXTIII.html>. Accessed 16 May 2025
55. Discontinued Product Ultra-High-Speed Modular Z:TA-R YSM40R Overview. <https://global.yamaha-motor.com/business/smt/mounter/ysm40r/>. Accessed 16 May 2025

56. SIPLACE SX – ASMPT SMT Solutions. <https://smt.asmpt.com/en/products/placement-solutions/siplace-sx/>. Accessed 16 May 2025
57. Hiller J, Lipson H (2007) Methods of Parallel Voxel Manipulation for 3D Digital Printing
58. 1.3.4. Broadside Coupling. <https://www.intel.com/content/www/us/en/docs/programmable/683624/current/broadside-coupling.html>. Accessed 19 May 2025
59. Han A, Zaderej V, Fitzpatrick R (2021) [Application Specific Electronics Package – Electronics Manufacturing Technology of the Future](#). In: 2021 14th International Congress Molded Interconnect Devices (MID). pp 1–6
60. Zaderej V, Fitzpatrick R, Arfaei B (2018) ASEP (APPLICATION SPECIFIC ELECTRONICS PACKAGE) A NEXT GENERATION ELECTRONICS MANUFACTURING TECHNOLOGY
61. (2020) ASEP (Application Specific Electronics Package) MID . 23(6):471–475. <https://doi.org/10.5104/jiep.23.471>
62. Tool Changer and Tools. <https://clank.tools/tools/>. Accessed 16 May 2025
63. (2017) What Are Molded Interconnect Devices or MIDs? <https://resources.altium.com/p/what-are-molded-interconnect-devices-or-mids>. Accessed 16 May 2025
64. Zhu W, Satterthwaite PF, Jastrzebska-Perfect P, Brenes R, Niroui F (2022) Nanoparticle contact printing with interfacial engineering for deterministic integration into functional structures. Science Advances 8(43):eabq4869. <https://doi.org/10.1126/sciadv.abq4869>
65. Zhu W, Satterthwaite PF, Niroui F (2022) [Deterministic and Scalable Patterning of Plasmonic Nanoparticles](#). In: 2022 Conference on Lasers and Electro-Optics (CLEO). pp 1–2
66. Han J, Niroui F, Lang JH, Bulović V (2022) Scalable Self-Limiting Dielectrophoretic Trapping for Site-Selective Assembly of Nanoparticles. Nano Lett 22(20):8258–8265. <https://doi.org/10.1021/acs.nanolett.2c02986>
67. Smith MM (2023) [Recursive Robotic Assemblers](#). Thesis, Massachusetts Institute of Technology
68. Blackburn LC (2021) [Superconducting Asynchronous Logic for Ultra-low Power High Performance Computing](#). Thesis, Massachusetts Institute of Technology

69. Blackburn LC, Wynn A, Berggren KK, Gershenfeld N (2025) A Compact Bit Serial Memory Cell for Adiabatic Quantum Flux Parametron Register Files. IEEE Trans Appl Supercond 35(5):1–5. <https://doi.org/10.1109/TASC.2025.3540048>