ECE8803-AI4

AI for Semiconductor Manufacturing

Module / Reinforcement Learning · Lecture 2

Reinforcement Learning for Plasma Etch

Asif Khan

Associate Professor
School of Electrical and Computer Engineering
School of Materials Science and Engineering
Georgia Institute of Technology

akhan40@gatech.edu

Course website https://aikhan123.github.io/khan-lab-website/ai-semiconductors-course.html

YouTube playlist https://www.youtube.com/watch?v=JOOhPoMm58I&list=PLF-ca21MPQt4

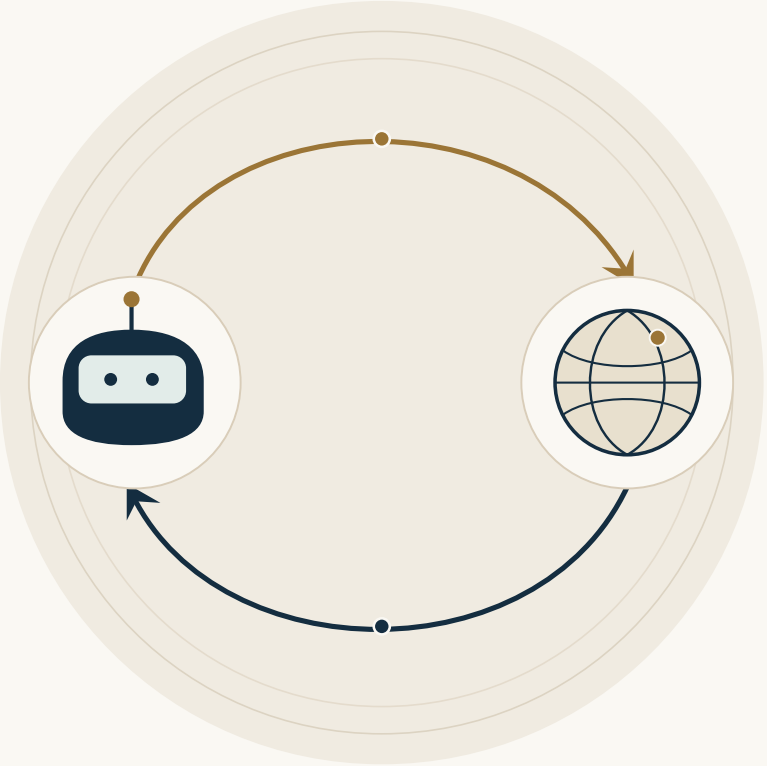

In this lecture

The Bosch deep silicon etch, step by step
A ViennaPS simulator of the trench
An agent that adjusts every cycle
Tests on new wafers and a second model

What this lecture does

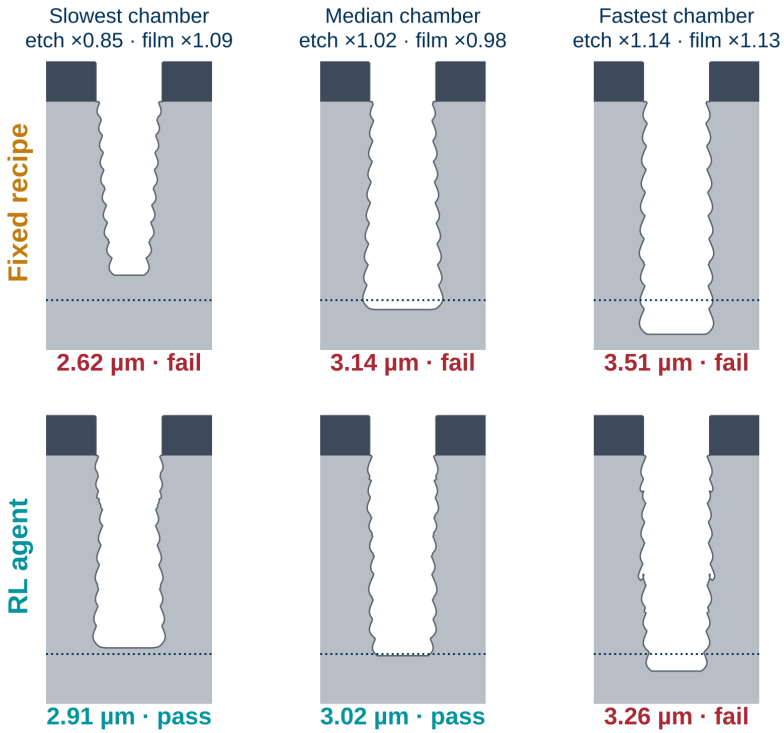

Slowest, median and fastest chamber among the 100 test wafers.
Dotted line: the 3.0 µm target.

1 The problem. The chamber drifts from wafer to wafer. A fixed recipe uses the same times anyway, so the depth follows the drift. 15% of wafers hit the target.

2 What we do. An RL agent looks at the trench after every cycle and picks the next cycle's film and etch times. It learns how from simulated wafers. 81% hit the target.

3 The check. An engineer's hand-written rule does the same job at 100%. So we also ask when RL is worth it.

From Connect Four to an etch chamber

	Connect Four (Lecture 1)	Bosch etch (today)
State	The board, with the player to move	The trench after each cycle: depth and width
Action	A column	Passivation time and etch time for the next cycle
Reward	+1 for a win, −1 for a loss, at the end	Minus the wall error each cycle, minus the depth error at the end
Episode	1 game, up to 42 moves	1 wafer, 10 cycles
Experience	Games of self-play	Runs of a ViennaPS process simulator
Learning	Average returns; Q-learning with ε-greedy	The same tabular Q-learning with ε-greedy

Same algorithm. The board becomes a trench cross-section.

When reinforcement learning works well

Lecture 1 ended with 3 conditions. Keep them in mind; we check each one at the end.

Decisions

Earlier actions change the options available later.

Feedback

The reward tracks the result we actually care about.

Experience

A simulator or safe trials provide many useful examples.

A Bosch etch runs 10 cycles in a row, ends in a measurable profile, and can be simulated. That makes it a good first test.

Deep silicon etching: narrow, deep, straight

Where it is used. MEMS sensors such as accelerometers, gyroscopes and microphones; through-silicon vias that connect stacked chips; deep trench capacitors and isolation; trench power transistors.

What matters. Walls that stay vertical, a bottom at the right depth, and the same result on every wafer.

Plain etching fails. Fluorine etches silicon at the same rate in every direction, so the trench widens as fast as it deepens.

Today's target. A trench 1.0 µm wide and 3.0 µm deep (aspect ratio 3), with vertical walls.

Production tools reach aspect ratios of 20 or more; 3 keeps each simulation under 1 s.

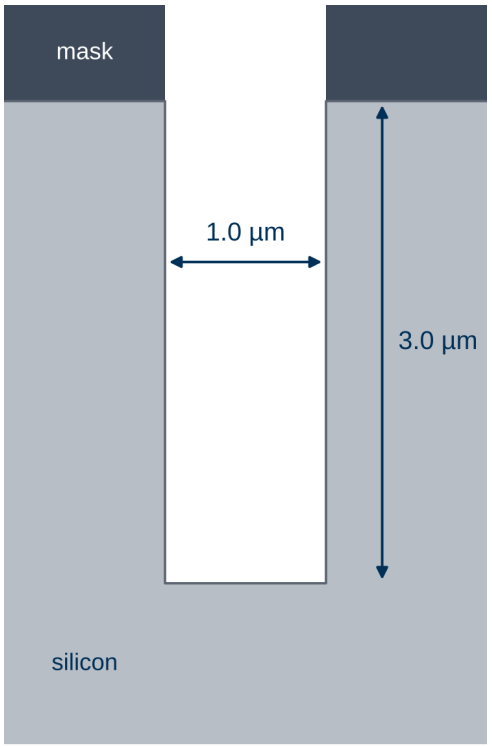

The Bosch process alternates 2 plasmas

Passivate C_4F_8 plasma

Deposits a thin fluorocarbon film on every surface.

Etch SF_6 plasma

Ions clear the film from the floor; fluorine etches the exposed silicon.

repeat 10 times

Franz Laermer and Andrea Schilp patented the idea at Bosch in the 1990s. It is still the standard way to etch deep silicon.

The trick. The film protects the walls, and the ions remove it only from the floor. Each etch step can only cut down.

The recipe. Gas, power and pressure for each step, the time of each step, and the number of cycles.

Today's knobs. Passivation time and etch time, chosen fresh for every cycle.

F. Laermer and A. Schilp, Robert Bosch GmbH, “Method of anisotropically etching silicon,” US Patent 5,501,893 (1996).

One cycle in 4 steps (ViennaPS)

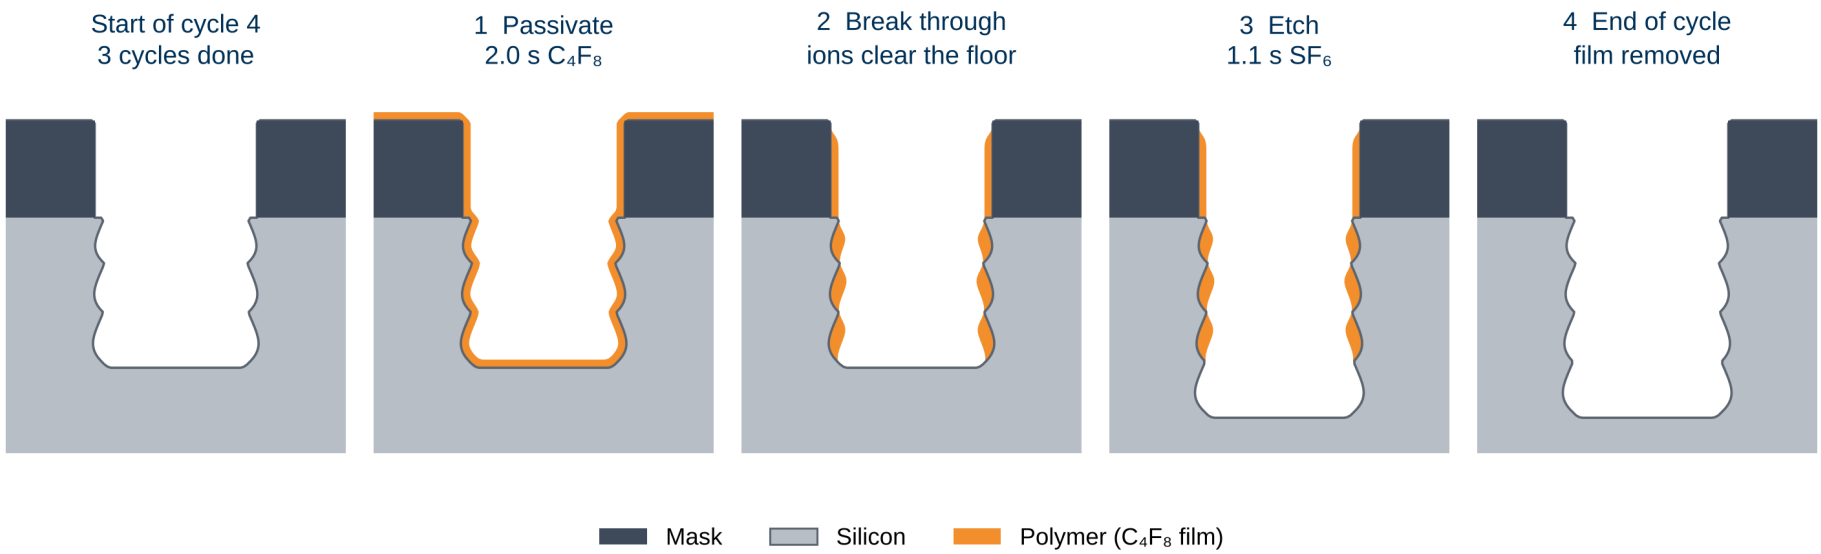

Each cycle cuts about 0.31 µm deeper and leaves one curved notch, a **scallop**, on each wall.

Simulated cross-section of cycle 4 with 2.0 s passivation and 1.1 s etch.

Why the film protects the walls but not the floor

Ions fall almost straight down. The wafer bias pulls them vertically, so they strike floors head-on and barely touch the walls.

Breakthrough. Ions remove 0.10 µm of film from horizontal surfaces. The film is 0.05 µm thick, so the floor clears and the walls keep their coat.

Etch. Fluorine etches the bare silicon: 0.20 µm/s downward with ion help, plus 0.08 µm/s in every direction.

The scallop. The sideways part cuts under the film. It stops when the step ends, and the next film covers the new notch.

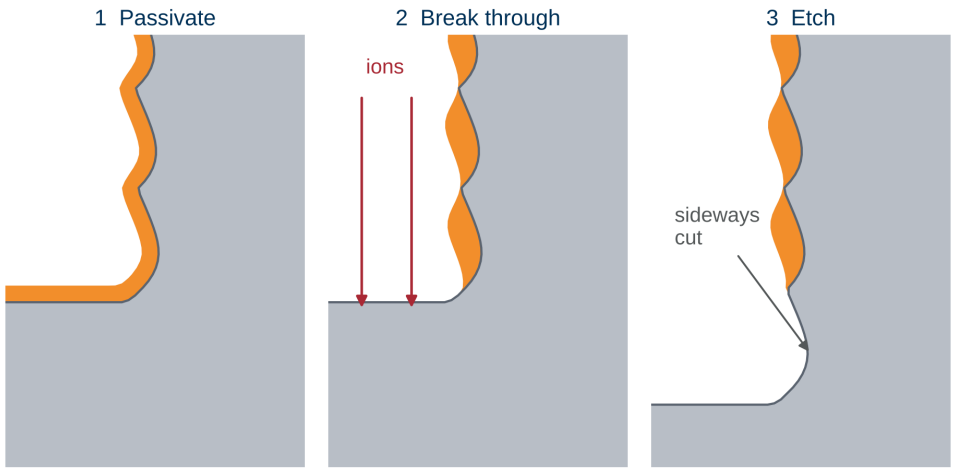

Close-up of the right wall in cycle 4

Repeating the cycle digs the trench

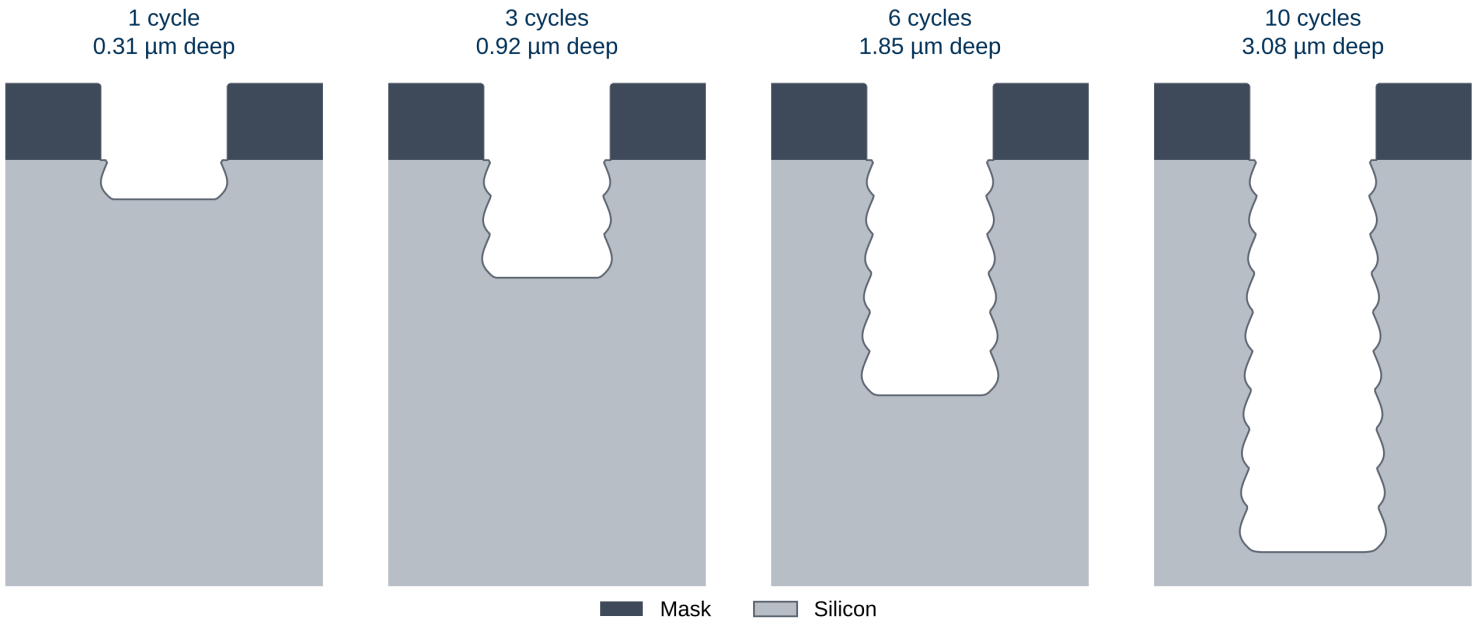

Same recipe each cycle: 0.31 µm per cycle, 3.08 µm after 10 cycles. The walls keep their scallops.

When the balance is off

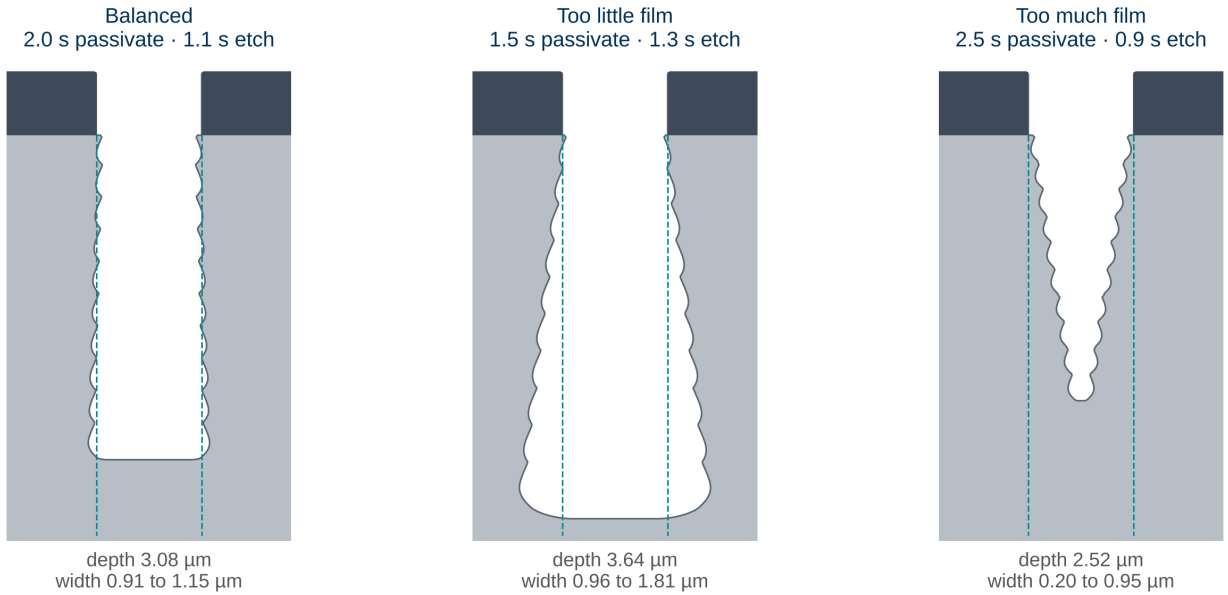

Too little film. The walls are etched through each cycle. The trench bows outward.

Too much film. The floor never fully clears. Each cycle starts narrower, so the trench tapers to a point.

Dashed lines mark the 1.0 µm mask opening.

Passivation and etch time have to match. The best match depends on how fast the chamber etches and deposits.

The recipe in numbers

Setting	Value	What it does
Passivation time	1.5, 2.0 or 2.5 s	Film grows 0.025 µm/s, so 2.0 s gives 0.05 µm.
Breakthrough	fixed	Ions remove 0.10 µm of film from horizontal surfaces.
Etch time	0.9, 1.1 or 1.3 s	Silicon etches 0.20 µm/s downward plus 0.08 µm/s in every direction.
Cycles	10	Target: 3.0 µm deep, so about 0.30 µm per cycle.
Mask	1.0 µm opening	0.6 µm thick. The mask itself is not etched.
Grid	0.02 µm	Spacing of the simulation grid.

A fixed recipe uses the same 2 times in every cycle. The agent may change them after each cycle.

Rates are model values chosen for teaching. In a tool, gas flows, RF powers, pressure and temperature set these rates.

The simulator: ViennaPS

ViennaPS is an open-source process simulator. It stores the wafer surface on a grid and moves it with a local etch or deposition speed at every surface point.

Geometric model (training)

We set the rates directly: downward, sideways, film growth.

10 cycles take about 0.35 s.

Fast enough for thousands of practice wafers.

Ray-traced model (check)

Ions and fluorine are traced from the plasma into the trench.

Rates follow the local particle flux, so deep floors receive less.

10 cycles take 1 to 2 minutes.

Train on the fast model. Test on the slower, more physical one.

T. Reiter, L. Filipovic et al., “ViennaPS: a process simulation framework,” SoftwareX (2025). Open source, TU Wien. Version 4.7 used here.

The spec and the score for 1 wafer

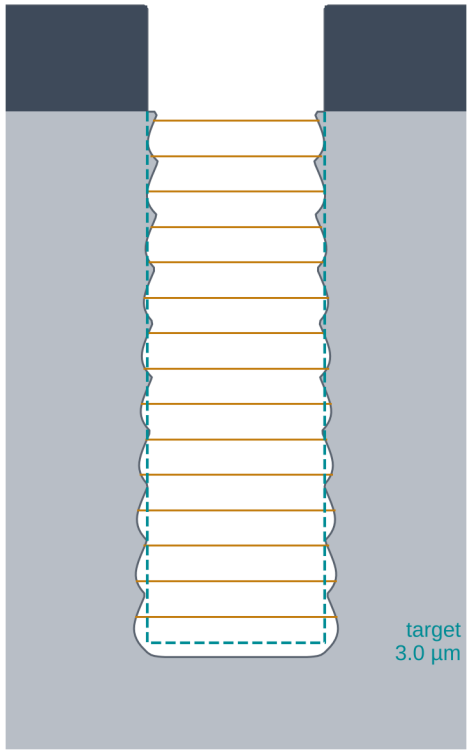

Measure. The width every 0.05 µm down the trench (60 lines here), and the depth.

Pass. Depth within 3.0 ± 0.15 µm, and every width within 1.0 ± 0.20 µm.

Score. Minus the sum of 2 errors:

depth error = |depth − 3.0 µm|

wall error = average |width − 1.0 µm|

This wafer. Depth 3.08 µm, so a depth error of 0.08 µm. Wall error 0.057 µm. Score −0.137. It passes.

Higher is better. A perfect trench scores 0.

No 2 wafers see the same chamber

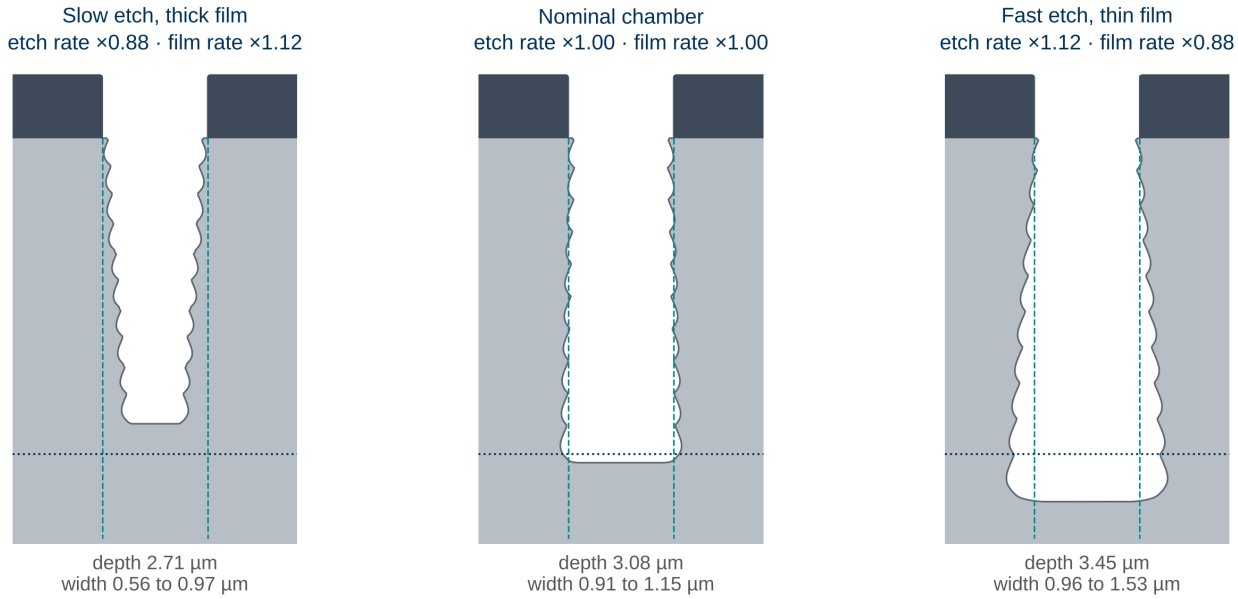

Chamber walls, gas purity, temperature and RF drift shift the rates from wafer to wafer.

Model. Etch rate × 0.85 to 1.15 and film rate × 0.85 to 1.15, drawn at random for each wafer.

The controller is never told these numbers.

Same recipe, 3 chambers: one tapers, one is on target, one bows and overshoots.

First, the best fixed recipe

Try all 9 combinations on 30 tuning wafers, each with its own random chamber. Same 2 times in every cycle.

Passivation time \ Etch time per cycle	0.9 s	1.1 s	1.3 s
1.5 s	score −0.53 pass 0%	score −0.53 pass 0%	score −1.16 pass 0%
2.0 s	score −0.57 pass 7%	**score −0.39** **pass 13%**	score −0.96 pass 3%
2.5 s	score −0.87 pass 0%	score −0.51 pass 7%	score −0.88 pass 0%

Best: 2.0 s and 1.1 s.

Average score −0.39

Passes on 13% of the tuning wafers.

This is the baseline the agent has to beat.

Even the best fixed recipe fails on most wafers.

A fixed recipe cannot follow the chamber

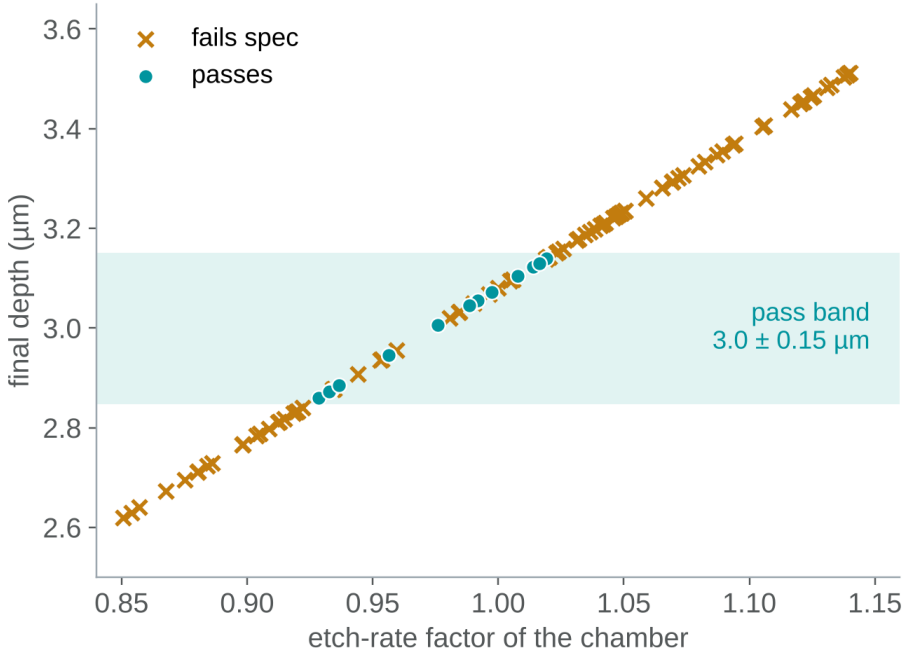

100 new test wafers, all run with 2.0 s and 1.1 s.

Depth follows the etch rate. A chamber that etches 10% faster digs a trench about 10% deeper, because the recipe never looks at the wafer.

32% land inside the depth band. 15% also keep every width within spec.

The film rate matters too: it sets whether that wafer bows or tapers.

To hit the target on every wafer, the recipe has to react to what it sees.

Feedback: measure after each cycle, then adjust

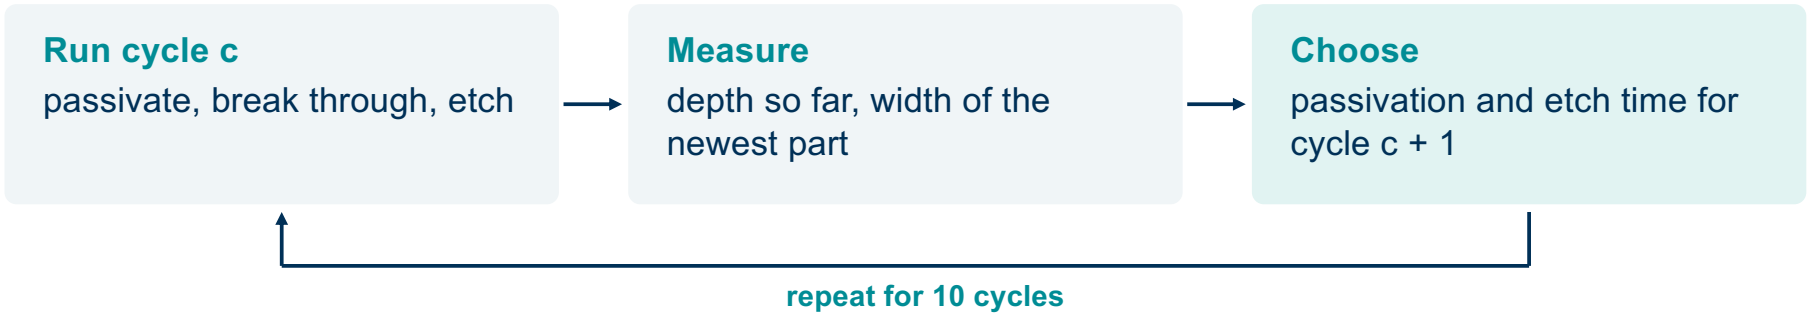

Fixed recipe = open loop. The same times every cycle, whatever happens.

Policy = closed loop. The next times depend on what was just measured.

Can a tool measure this? Depth: yes, laser interferometry tracks it during the etch.

Width: not directly; it needs a model or a sensor proxy. Here the simulator reports both.

The etch problem as a Markov decision process

Ingredient	Meaning	In the etch
States **S**	What the controller knows before each cycle.	Cycle number, depth vs. schedule, newest width: 90 states.
Actions **A(s)**	The choices available in state s.	9 pairs of passivation time and etch time.
Transitions **P(s′ | s, a)**	How the next state follows.	1 ViennaPS cycle with this wafer's hidden rates.
Reward **R(s, a, s′)**	The reward for that step.	Minus the wall error the cycle added. After cycle 10, also minus the depth error.
Discount **γ**	Weight on later rewards.	γ = 1. Every wafer ends after 10 cycles.
Markov property	**Only roughly. The hidden chamber rates carry over from cycle to cycle; the readings reveal them a little at a time.**	

The state: 3 coarse readings

Reading	Values
Cycle	0 to 9: how many cycles are done.
Depth vs. schedule	Behind, on, or ahead of 0.30 µm per cycle (±0.06 µm).
Newest width	Narrow (< 0.96 µm), on target, or wide (> 1.06 µm), over the last cycle's cut.

Example. A fast chamber after 4 cycles: 1.38 µm deep against a 1.20 µm schedule, newest width 1.19 µm.

State = (4, ahead, wide).

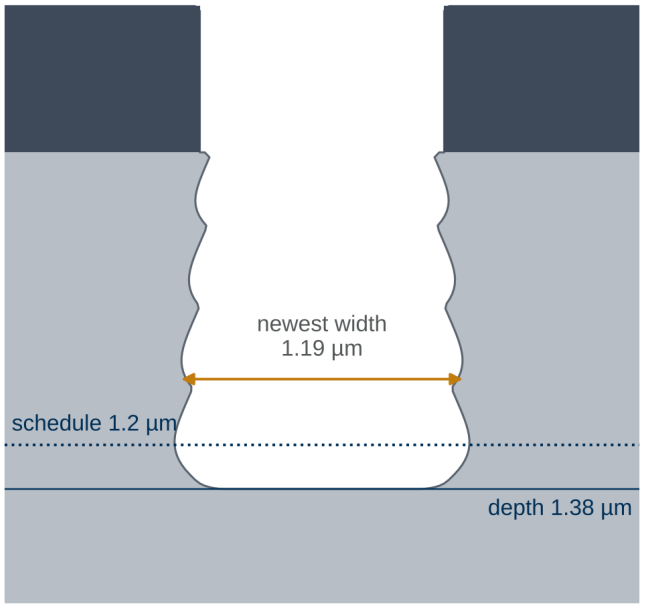

Fast etch, thin film, after 4 cycles of the fixed recipe

The actions: 9 recipe choices per cycle

Passivation time \ Etch time	0.9 s	1.1 s	1.3 s
1.5 s	**action 0** film 0.038 µm cut 0.25 µm	**action 1** film 0.038 µm cut 0.31 µm	**action 2** film 0.038 µm cut 0.36 µm
2.0 s	**action 3** film 0.050 µm cut 0.25 µm	**action 4** film 0.050 µm cut 0.31 µm	**action 5** film 0.050 µm cut 0.36 µm
2.5 s	**action 6** film 0.062 µm cut 0.25 µm	**action 7** film 0.062 µm cut 0.31 µm	**action 8** film 0.062 µm cut 0.36 µm

Longer etch: deeper, and a deeper scallop.

Longer passivation: thicker film, so less sideways cutting; too thick and the floor stays covered.

Film thickness and depth cut per cycle on a nominal chamber.

Choices are kept inside a window around the fixed recipe, as an engineer would allow on a qualified tool.

The reward has to count what we care about

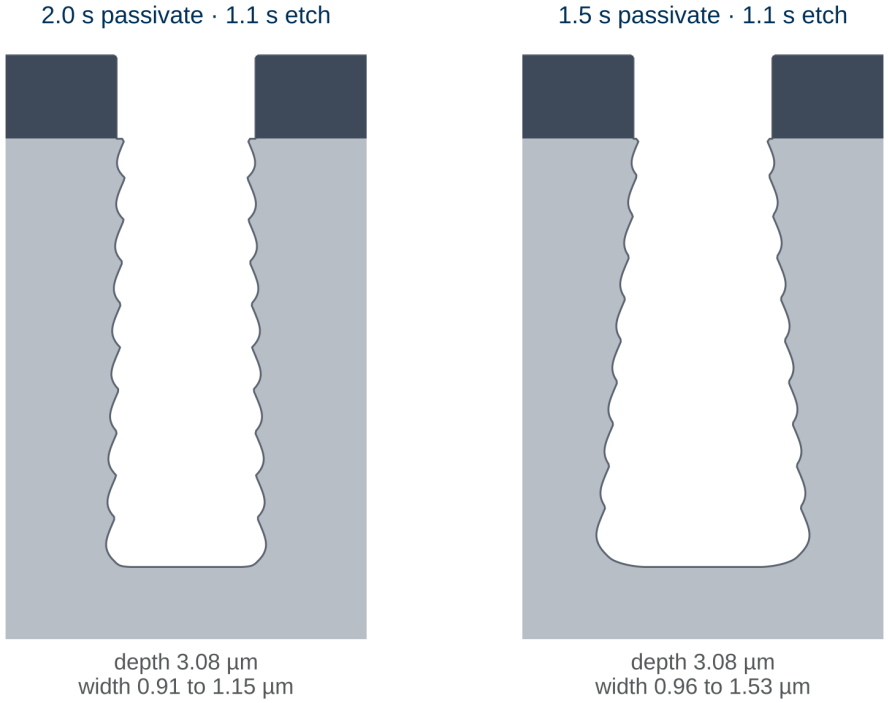

Our reward. After each cycle: minus the wall error that cycle added. After cycle 10: also minus the depth error. The rewards add up to the wafer's score.

A reward that only counts depth gives these 2 wafers the same marks: both are 3.08 µm deep.

The right one bows to 1.53 µm. An agent paid for depth alone has no reason to avoid it.

Per-cycle rewards also tell the agent early which cycle went wrong.

Learning: Q-learning, as in Lecture 1

Table	Q(s, a) for 90 states × 9 actions, all starting at 0.
Episode	1 simulated wafer: 10 cycles, a fresh random chamber each time.
Choose	ε-greedy: a random action with probability ε, else the best Q. ε falls from 0.30 to 0.02.
Update	After each cycle, move Q(s, a) toward the target: this cycle's reward + the best Q of the next state.
Budget	4,000 wafers, about 33 minutes on 1 CPU core.

Same rule as the Connect Four table: new Q = old Q + step × (target − old Q).
Only the state, actions and reward changed.

Training: the score climbs, then levels off

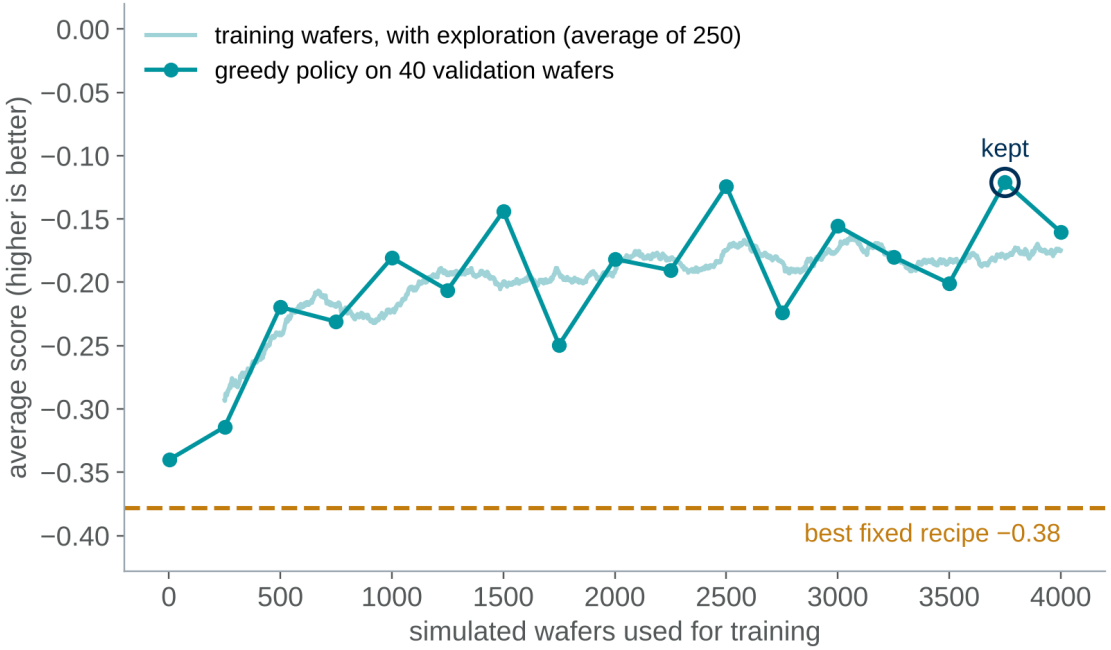

Start. After 1 training wafer the table is almost empty. Score −0.34, near the fixed recipe.

After training. The greedy policy scores about −0.12 on the validation wafers; the fixed recipe scores −0.38 on the same wafers.

Noisy. Q estimates are averages of noisy wafers, so the greedy choice flips between close actions.

Kept the table with the best validation score (wafer 3,750). The 100 test wafers are used only once, at the end.

Most of the gain comes in the first 500 wafers. After that the greedy policy wobbles between −0.25 and −0.12.

What the agent learned

Average times the policy chose on the test wafers, by reading (cycles 2 to 10).

Depth vs. schedule	Newest width: narrow	Newest width: on target	Newest width: wide
behind	**1.6 s · 1.30 s** 48 visits	**2.0 s · 1.27 s** 51 visits	**2.2 s · 1.28 s** 19 visits
on	**1.7 s · 1.19 s** 78 visits	**2.1 s · 1.05 s** 221 visits	**2.2 s · 1.01 s** 63 visits
ahead	**1.8 s · 0.91 s** 40 visits	**2.0 s · 1.01 s** 257 visits	**2.5 s · 0.92 s** 123 visits

passivation time · etch time

Etch time follows depth.
Behind schedule: about 1.3 s. Ahead: 0.9 to 1.0 s.

Passivation follows width.
Narrow: 1.6 to 1.8 s. Wide: 2.2 to 2.5 s.

Nobody wrote these rules. They came from rewards alone.

2 feedback rules an etch engineer would recognize, learned from simulated wafers.

One wafer, cycle by cycle

Fixed recipe

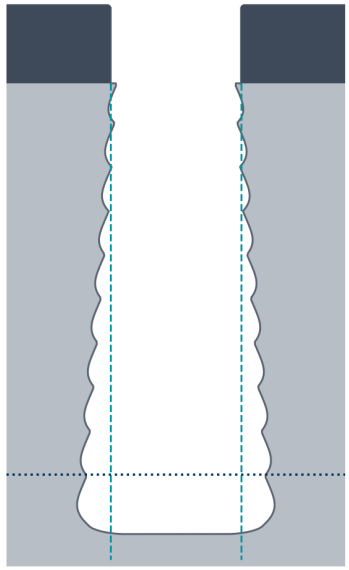

depth 3.45 µm
width 0.96 to 1.52 µm
score −0.70 · fail

Learned policy

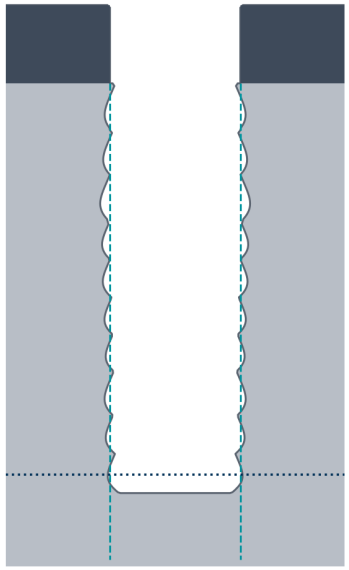

depth 3.14 µm
width 0.94 to 1.15 µm
score −0.20 · pass

Fast etch, thin film: etch rate × 1.12, film rate × 0.90

Cycle	Reading before	Policy chose
1	start	1.5 s · 1.3 s
2	ahead, ok	2.5 s · 1.1 s
3	ahead, wide	2.0 s · 1.1 s
4	ahead, wide	2.5 s · 0.9 s
5	ahead, wide	2.5 s · 0.9 s
6	ahead, wide	2.5 s · 0.9 s
7	ahead, wide	2.5 s · 0.9 s
8	ahead, ok	2.5 s · 1.1 s
9	ahead, wide	2.5 s · 0.9 s
10	ahead, ok	2.5 s · 0.9 s

From cycle 2 on, the trench runs ahead, so the policy thickens the film and shortens the etch.

Test on 100 new wafers

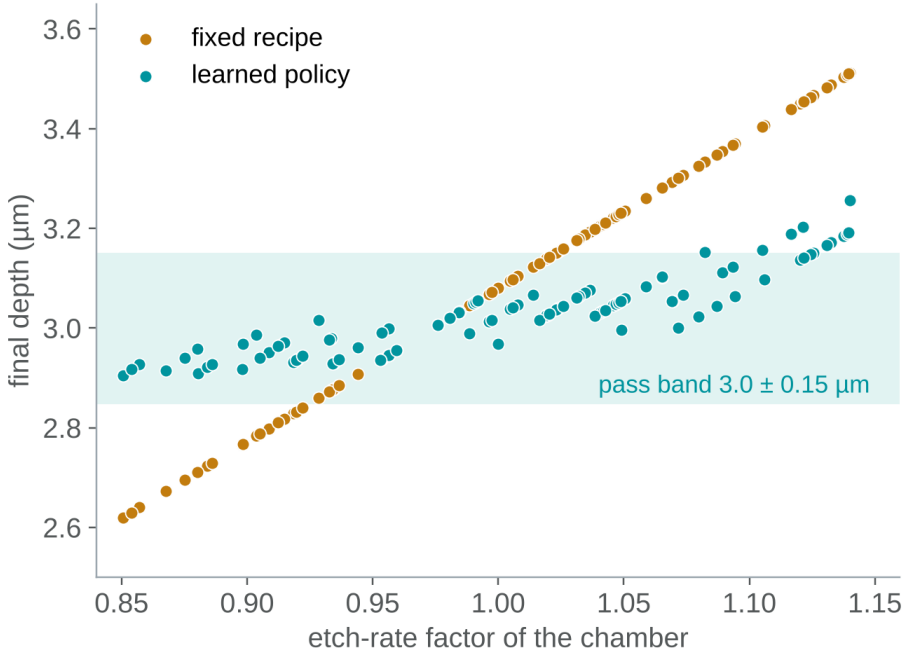

	Fixed recipe	Policy
Pass rate	15%	81%
Average score	−0.35	−0.13
Depth spread (1σ)	0.25 µm	0.08 µm
Wall error (avg)	0.118 µm	0.058 µm
Worst wafer score	−0.79	−0.36

The policy scores higher on 97 of the 100 wafers.

Depth no longer follows the etch rate: the spread drops from 0.25 to 0.08 µm.

Pass rate 15% → 81% on wafers the agent never saw.

Check on a second model: ray-traced ViennaPS

Fixed recipe

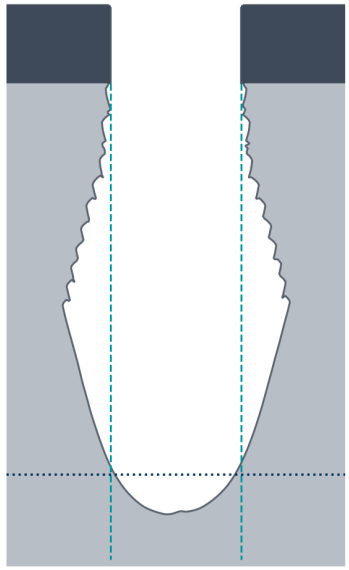

depth 3.30 µm
width 0.52 to 1.75 µm
score −0.66 · fail

Learned policy

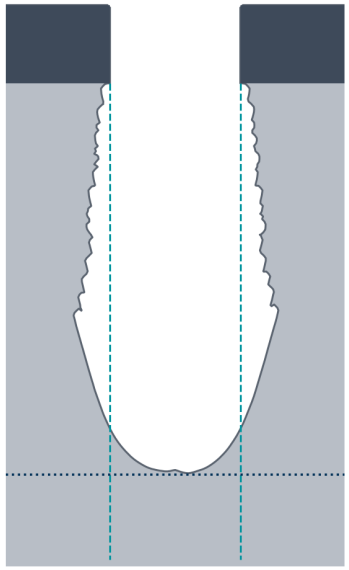

depth 2.99 µm
width 0.66 to 1.58 µm
score −0.32 · fail

Same 8 test wafers, run on the ray-traced model. The policy was trained only on the geometric model.

	Fixed recipe	Policy
Pass rate	0 of 8	0 of 8
Average score	−0.36	−0.29
Depth range	2.66–3.30 µm	2.88–3.01 µm

New effects: the etch slows as the trench deepens, and the lower walls bow.

Depth transferred. Every policy wafer ends within 0.12 µm of 3.0 µm.

Walls did not. The rounded floor reads narrow, so the policy thins the film. Wall error 0.19 → 0.23 µm.

A policy inherits its simulator's blind spots. Fix the model, or retrain on the better one, before trusting either.

From the simulator to a real chamber

Calibrate
Fit the simulator's rates to measured cross-sections (SEM) from the real tool.

Randomize
Train over a range of chambers, as we did with the hidden rate factors, so the policy does not rely on one.

Measure
Track depth in situ with interferometry; estimate width with a virtual-metrology model.

Fence
Allow only actions inside a qualified window. Fall back to the fixed recipe if a reading looks wrong.

Shadow
Let the agent suggest while the engineer decides. Compare on test wafers before release.

Update slowly
Learn from real wafers in small steps, with a human sign-off.

Each step closes part of the gap we just saw between 2 models.

Machine learning for etch recipes in industry

Hitachi: reinforcement learning for multi-step recipes

Etch recipes now have many steps, each with its own plasma settings.

RL chose the settings step by step, to reach a deep, vertical trench.

A 5-step recipe met the target profile within 100 experimental trials.

Lam Research: humans and algorithms on a virtual etch

A profile simulator stood in for the lab. Engineers and Bayesian optimization competed to hit a target at lowest cost.

Engineers were best early; the algorithm was best near the tight target.

Human first, computer last cut the cost to target by half.

Both used a simulator or few real trials. Both kept engineers in the loop.

Bayesian optimization (Module 9) picks 1 recipe; reinforcement learning picks a recipe for each step.

Nakada, Usui, Ishikawa, Ohmori (Hitachi), APS GEC 2018 and ICDDPS 2019. Kanarik et al. (Lam Research), Nature 616, 707–711 (2023).

Checking the 3 conditions

Condition	In the Bosch etch	Verdict
Decisions	Each cycle's times change the trench the next cycle starts from.	Met: 10 linked choices per wafer.
Feedback	Per-cycle wall error and final depth error add up to the score we care about.	Met, if the tool can measure depth and width.
Experience	4,000 simulated wafers cost minutes, not a fab's budget.	Met in simulation. The model check shows the remaining gap.

All 3 hold in simulation. The open question is the third: how well the simulator matches the real tool.

Key ideas from this lecture

The process	Bosch etching alternates film and etch. Their balance sets the wall shape.	5–9
The problem	Chamber rates drift from wafer to wafer, and a fixed recipe cannot follow them.	13–15
The MDP	State: cycle, depth vs. schedule, newest width. Action: 2 times. Reward: wall and depth error.	17–20
The learning	Same tabular Q-learning and ε-greedy as Connect Four, on simulated wafers.	21–22
The result	On 100 new wafers the pass rate rose from 15% to 81%. Etch time follows depth; passivation follows width.	23–25
The caution	A policy is only as good as its simulator. Check on a second model, then fence and shadow on the tool.	26–27

Lab: train your own agent on this ViennaPS environment, then change the state, the reward, or the chamber range.