



Department of Precision and Microsystems Engineering

Wafer-Based Thin-Film Circular Cutting Edges at High Speed

Thomas Schelling

| | |
|----------------|------------------|
| Report no | : 2026.024 |
| Professor | : Dr. R.A. Norte |
| Specialisation | : MSD |
| Type of report | : MSc Thesis |
| Date | : 16 March 2026 |

WAFER-BASED THIN-FILM CIRCULAR CUTTING EDGES AT HIGH SPEED

MSc Thesis

by

Thomas SCHELLING

Student Number: 4963407
Thesis advisor: Dr. R.A. Norte, TU Delft, ME faculty
Date: 16 March 2026



ABSTRACT

This thesis investigates whether a wafer-defined circular blade can realise a mechanically robust cutting edge at thin-film thickness scale and perform first material removal under rotation. The concept uses a silicon annulus carrying a deposited film. Selective substrate removal near the outer perimeter releases a circumferential overhang whose thickness defines the nominal kerf scale while overhang length governs achievable cutting depth. The work is structured around four challenges: reaching the minimum cutting-edge speed with acceptable run-out, fabricating annular blade blanks, selectively underetching the rim while preserving mounting geometry, and integrating the system to demonstrate cutting.

A practical spindle operating envelope was established in which the blade reached the minimum target speed while operational displacement remained predominantly rotation-synchronous and below the realised overhang range used for survival and engagement trials. LS-Precise femtosecond-laser cutting enabled repeatable fabrication of 525 μm -thick silicon annular blanks, and a lid-base sacrificial masking concept enabled selective SF_6 underetching and controlled release of SiO_2 rims while preserving hub-bore integrity when the lid remained intact. In the integrated demonstration, released SiO_2 rims survived operation at and above target speed and produced a continuous kerf in rigid ABS under incremental approach at $N = 17.7\text{ kRPM}$, with measured trench depths of 11.72 μm to 14.71 μm and a central kerf width of $\sim 46\ \mu\text{m}$. The main limitations are that rigid cutting was attributable primarily to the rear-side $\sim 10\ \mu\text{m}$ rim, sustained controlled engagement of the front-side $\sim 4\ \mu\text{m}$ ultra-thin rim remains unproven, and circumferentially continuous ultra-thin a-SiC edges were not reliably achieved within the present process window.

ACKNOWLEDGEMENTS

This thesis has been quite a journey. What began with an idea inspired by a YouTube video grew into a project that felt genuinely exploratory, and I greatly enjoyed that it allowed me to combine my interest in mechanical engineering with the challenge of designing, building, and integrating many different parts into one working system. Along the way, I was fortunate to receive the support of many people.

I would first like to thank Professor Richard Norte for his enthusiasm for this project and for the trust he placed in both the work and in me personally. That trust meant a great deal, especially during periods when I doubted myself or the progress of the project. I also greatly appreciated his openness and the relaxed, honest way in which we could communicate. It always felt possible to speak freely, not only about the project itself, but also about how I was doing more broadly.

I am very grateful to Paulina Castro and Lucas Norder for always taking the time to look for and provide me with the materials I needed for this project. I would like to thank Lucas in particular for carrying out the etching process. I would also like to thank them both for their advice and for thinking along with the project as it developed. My thanks also go to everyone in the Norte lab team for their interest, encouragement, and willingness to ask questions and help think through problems. I especially appreciated the collective enthusiasm at the moment when the blade first performed a successful test.

I would furthermore like to thank the PME lab technicians for their patience with the many questions and requests I brought to them over the course of this thesis. Their help was invaluable, and I always appreciated how readily they engaged with the technical challenges and helped think through possible solutions. I am equally thankful to the faculty workshop team, and in particular to Sam Berlemon and Damian de Nijs, for their support and expertise. Without their help, I would not have been able to produce two very high-precision parts that were essential to this work.

Finally, I would like to thank my family for their unwavering support and trust in me throughout this project. Their encouragement carried me through the difficult stages, and, in a very practical sense, this setup would also not have come together without their help, including the expertise they generously shared on dry ice handling and bearing assembly techniques.

Thomas Schelling
Amsterdam, February 2026

CONTENTS

| | |
|---|------------|
| Abstract | iii |
| Acknowledgements | v |
| List of Figures | ix |
| List of Tables | xv |
| 1 Introduction | 1 |
| 2 Challenge 1: Rotation to cutting-edge speed with acceptable run-out | 5 |
| 2.1 Practical challenges in achieving micrometre-level run-out at high rotational speeds. | 6 |
| 2.2 Design approach and baseline configuration | 6 |
| 2.2.1 Balancing strategy | 7 |
| 2.3 Bearing selection and axial constraint concept | 7 |
| 2.4 Manufacturing of bearing seats and fit verification | 8 |
| 2.5 Verification of speed and displacement measurements | 9 |
| 2.5.1 Tachometry: speed measurement verification | 9 |
| 2.5.2 Displacement measurement verification. | 10 |
| 2.6 Setup performance: quasi-static run-out and operational displacement versus speed | 11 |
| 2.6.1 Test matrix and measurement points | 12 |
| 2.6.2 Effect of an assembly-induced axial bias load | 12 |
| 2.6.3 Quasi-static run-out and assembly sensitivity (TIR) | 12 |
| 2.6.4 Operational displacement versus rotational speed (RSR and WCR) | 13 |
| 2.6.5 Acceptance versus overhang length | 13 |
| 2.7 Discussion and implications | 14 |
| 3 Challenge 2: Femtosecond-laser fabrication of annular blade blanks | 17 |
| 3.1 Constraint-driven blade outline selection. | 17 |
| 3.2 LS-Precuss enables through-cuts in 525 μm silicon | 17 |
| 3.3 Production recipe, reliability–time trade-off, and yield | 18 |
| 3.4 Edge condition and implication for minimum underetch length | 18 |
| 3.5 Bore-fit tuning under a handling constraint. | 19 |
| 3.6 Discussion and implications | 20 |
| 4 Challenge 3: Selective underetch with hub-bore protection | 21 |
| 4.1 Hub-bore protection during plasma etching | 22 |
| 4.1.1 Lid fabrication | 23 |
| 4.2 Blade fixturing and handling for plasma etching | 23 |
| 4.2.1 Base fabrication and contact-surface topography | 25 |
| 4.2.2 Selection of target undercut (overhang) length. | 25 |
| 4.3 Etching results and fixture performance | 26 |
| 4.3.1 Summary of etch trials | 26 |
| 4.3.2 Etching results for Si/a-SiC blades | 27 |
| 4.3.3 Etching results for Si/SiO ₂ blades and hub-bore protection | 28 |
| 4.3.4 Released-edge curvature and stress-gradient consistency check (SiO ₂) | 29 |
| 4.4 Discussion and implications | 30 |
| 5 Challenge 4: Integrated proof-of-concept cut demonstration | 33 |
| 5.1 Experimental constraints and approach | 33 |
| 5.2 Survival and compliant engagement tests. | 33 |
| 5.2.1 Test 1: Survival at speed | 33 |

| | | |
|----------|---|-----------|
| 5.2.2 | Tests 2–4: Cutting of a tensioned thread | 34 |
| 5.3 | Rigid-sample tests | 34 |
| 5.3.1 | Test 5: Over-travel and substrate sidewall contact | 35 |
| 5.3.2 | Test 6: Controlled incremental rigid cut | 35 |
| 5.4 | Discussion and implications | 37 |
| 6 | Conclusion and outlook | 39 |
| 6.1 | Challenge-wise conclusions. | 40 |
| 6.2 | Outlook | 40 |
| 6.2.1 | First-contact detection, optical access, and controlled depth exploration | 40 |
| 6.2.2 | Ultra-thin rim robustness, attribution, and gradual failure concepts. | 41 |
| 6.2.3 | Spindle repeatability, operational displacement, and motor-drive effects | 41 |
| 6.2.4 | Plasma-etch process window and translation to a-SiC | 41 |
| 6.2.5 | Broader scope: wafer-defined mechanical edges as a fabrication platform. | 41 |
| A | Signal processing and robustness checks | 43 |
| A.1 | Order-domain displacement reconstruction from LDV velocity. | 43 |
| A.2 | Operating point definition under RPM drift. | 43 |
| A.3 | Bootstrap procedure for uncertainty estimates | 44 |
| A.4 | Selection of reconstruction order M_{recon} | 44 |
| A.5 | Harmonic distortion ratio (HDR) and interpretation | 44 |
| B | Supplementary results for Challenge 1 | 47 |
| B.1 | Baseline operational displacement at the shaft axis ($r = 0$ mm) | 47 |
| B.2 | Full axial blade-radii results (RSR and WCR_{95}) | 47 |
| B.2.1 | Order-domain decomposition across blade radii. | 47 |
| C | Supplementary figures for Challenge 2 | 51 |
| C.1 | Qualitative influence of laser power on near-edge morphology in LS-Precess mode | 51 |
| C.2 | Manual mounting limitation when using a vacuum tweezer | 51 |
| D | Initial bending and rotational stiffening | 53 |
| D.1 | Theoretical background: initial deflection | 53 |
| D.2 | Estimating the initial tip deflection of the cutting edge | 55 |
| D.2.1 | Initial bending estimates for the SiO_2 edge (comparison case). | 56 |
| D.3 | Rotational stiffening model | 57 |
| D.4 | Numerical evaluation of rotational stiffening | 57 |
| D.4.1 | Evaluation for the a-SiC design case | 57 |
| D.4.2 | SiO_2 comparison case (numerical anchor) | 58 |
| E | Rayleigh–Ritz derivation | 59 |
| E.1 | Assumptions | 59 |
| E.2 | Centrifugal force distribution | 60 |
| E.3 | Representing the initial curve | 60 |
| E.4 | Rayleigh–Ritz energy balance | 61 |
| E.5 | Evaluation of the integrals I_1 and I_2 | 61 |
| E.5.1 | Integral I_1 | 61 |
| E.5.2 | Integral I_2 | 62 |
| E.5.3 | Final expression | 62 |
| F | MATLAB code | 63 |
| G | Supplementary figures for Challenge 4 | 67 |
| G.1 | Imaging trials for clearance indication at the rotating rim. | 67 |
| G.2 | Sample holder for compliant engagement (thread tests) | 67 |
| G.3 | Sewing-thread response used as a qualitative contact indicator. | 68 |
| G.4 | Post-compliant sample test microscopy and interpretation. | 68 |
| G.5 | Post-test 5 microscopy and interpretation | 69 |
| G.6 | Representative WLI map and line profile (Test 6 kerf depth) | 69 |
| G.7 | Post-test 6 microscopy and interpretation | 69 |

LIST OF FIGURES

| | | |
|-----|--|----|
| 1.1 | Fabrication concept for the wafer-based circular blade (schematic, not to scale). A thin film on a silicon substrate is shaped into an annular disk, after which selective substrate removal near the outer edge releases a circumferential thin-film overhang that forms a continuous cutting edge. The highlighted region indicates the substrate removal zone. | 3 |
| 2.1 | Mechanical layout of the high-speed rotation test setup used in Challenge 1. A brushless DC motor drives a stainless steel shaft via a flexible jaw coupling; the shaft is supported by two deep-groove ball bearings in a rigid cylindrical housing mounted to an aluminium base plate. The inset section view highlights the locating/non-locating bearing concept (axial constraint by one bearing and axial float by the other) and the blade clamp arrangement, with the blade positioned close to the outboard bearing to minimise unsupported overhang. | 7 |
| 2.2 | Integrated sensing configuration for tachometry and LDV measurements. An optical-reflective tachometer sensor is mounted above the 8 mm-diameter shaft (inset) to provide a pulse train used for RPM estimation and synchronous triggering (CH1). A Laser Doppler Vibrometer (LDV) measures <i>axial</i> vibration velocity and operational displacement is reconstructed from the velocity signal in post-processing. A single-axis translation stage repositions the LDV spot radially to select the measurement radius r (e.g. $r = 0$ on the shaft centre and $r = 15.5$ mm to 17.5 mm on the blade). The analogue decoder output is recorded on the oscilloscope (CH2). | 11 |
| 2.3 | Commissioning of the tachometer trigger signal on oscilloscope CH1. (a) With AC coupling enabled, the square-wave pulse train exhibits baseline droop and sloped plateaus (insets), consistent with attenuation of low-frequency content by the AC-coupling network. A fixed trigger threshold can therefore shift in time relative to the true edge. (b) With DC coupling, the waveform preserves stable logic levels and a constant baseline (insets), enabling reliable threshold triggering. In subsequent measurements, $N_{ppr} = 2$ pulses per revolution were generated and a $1\times$ reference was obtained by triggering on every second rising edge. | 11 |
| 2.4 | Operational displacement versus rotational speed at blade radii, showing the repeatable rotation-synchronous component (RSR) and per-revolution severity metric (WCR_{95}). (a) Axial LDV configuration at $r = 16.0$ mm and $r = 17.5$ mm. (b) Radial LDV configuration at the rim ($r \approx 18$ mm). Solid curves denote RSR; dashed curves denote WCR_{95} . Each marker corresponds to one 7 s time record and is plotted at its measured median RPM; n in the legend indicates the number of records in each series. | 15 |
| 3.1 | Comparison of femtosecond-laser through-cutting of 525 μm -thick silicon in normal and LS-Precess modes. (a) Normal mode: on-axis camera views of a power—repetition matrix (power 50 % to 80 %; repetitions 4000 to 8000). The photograph of the resulting die indicates no released features (no through-cuts) within this parameter window. (b) LS-Precess: matrix (power 60 % to 80 %; repetitions 3500 to 4500) with die and released-disk photographs confirming through-cuts at all tested points. Scan speed (200 mm/s), outlines (three), and outline pitch (0.01 mm) were constant. | 18 |
| 3.2 | Representative edge-condition observations after femtosecond-laser cutting, illustrating near-edge damage and a critical failure mode. (a) Si/SiO ₂ with the film facing the laser: a visibly affected near-edge region is present; annotated widths indicate representative local values (extent varies along the circumference). (b) Si/SiO ₂ with the film facing away from the laser: the visibly affected region is reduced at the shown location, indicating a modest orientation benefit. (c) Example failure mode after an incomplete through-cut: backside chipping and a local bulge at the kerf exit, where partially attached material subsequently broke away. Scale bars as indicated. | 19 |

| | | |
|-----|---|----|
| 3.3 | Supporting kerf-width and dimensional-bias characterisation for femtosecond-laser cutting in LS-Precess mode using simple test patterns. The finite kerf width causes the realised edges to deviate systematically from the programmed centreline, motivating empirical tuning of the programmed bore radius by stepwise fit tests on the shaft seat. Scale bars as indicated. | 20 |
| 4.1 | Schematic illustration of two geometric contributions to quasi-static radial run-out at the blade outer perimeter. Solid and dashed outlines indicate two opposite angular positions (180° apart). (a) A diametral clearance s between the shaft blade seat and the blade hub bore permits eccentric positioning of the blade relative to the shaft by up to $s/2$, which can contribute up to s peak-to-peak run-out at the outer perimeter. (b) Eccentricity between the centres of the hub bore and the outer perimeter by d introduces a geometric run-out of $2d$ peak-to-peak at the outer perimeter. | 21 |
| 4.2 | Lid concept for selective protection of the hub bore during isotropic SF_6 plasma etching. (a) Lid geometry: a stepped Si disk with outer radius 15 mm and a central locating tower (radius 8.5 mm) providing an approximately $180\ \mu\text{m}$ step height. (b) Cross-sectional schematic of the lid placed on the blade with the tower inserted into the hub bore, forming a self-aligned sacrificial hard mask that blocks plasma exposure of the hub bore sidewalls while leaving the outer perimeter exposed for substrate undercutting; the lid is held in place by gravity and can be handled using a vacuum tweezer. Arrows indicate incoming etch species and red shading indicates etch-exposed Si surfaces. | 23 |
| 4.3 | As-fabricated Si lid produced by femtosecond-laser through-cutting and depth-controlled pocket milling. (a) Lid manipulated using a vacuum tweezer, enabling handling without contacting the outer edge. (b) Close-up view showing the raised central locating tower (radius 8.5 mm) and the laser-machined contact face (outer radius 15 mm) that interfaces with the blade during etching (see inset for orientation). The tower step height was verified by white-light interferometry ($\approx 180\ \mu\text{m}$). | 23 |
| 4.4 | Base-assisted handling workflow for SF_6 plasma etching of the annular blade. (a) Base oil-bonded to carrier wafer (oil only at the base-carrier interface). (b) The blade is placed onto the base with the tower inserted into the hub bore, providing self-aligned lateral constraint while the blade is held by gravity. (c) Lid added to shield the hub-bore sidewalls while leaving the outer perimeter exposed (red: etch-exposed Si). (d) Photograph of the base positioned on the carrier wafer. (e) Photograph of the assembled stack (carrier wafer + base beneath, blade, and lid on top) prior to etching. | 24 |
| 4.5 | Centrifugally stiffened cantilever-beam model of the ultra-thin cutting edge. (a) Schematic of the rotating dicing blade, indicating the inner and outer blade radii r_i and r_o , the clamping radius R_E of the overhanging cutting edge, and the direction of rotation Ω . (b) The cutting edge is idealised as a slender cantilever of length L , in-plane width w , and thickness d , attached to the blade body at radius R_E . Local beam coordinate system with axial coordinate x measured from the clamped end ($x = 0$) to the free tip ($x = L$) and transverse deflection $w(x)$ in the z -direction, as used in the Rayleigh-Ritz analysis of centrifugal stiffening. | 25 |
| 4.6 | Optical microscopy of a Si/a-SiC blade before and after Etch 1 (SF_6 , 5 min; target undercut $L \approx 60\ \mu\text{m}$). (a) Pre-etch a-SiC surface. (b) Post-etch blade showing a matte-grey exposed annulus outside the lid footprint. (c) Close-up of the protected-exposed boundary demonstrating a sharp lid-defined transition. (d) Rough Si-like morphology in the exposed region consistent with a-SiC film loss. Red markers indicate magnified locations. Scale bars: 1 mm (a–b), $100\ \mu\text{m}$ (c–d). | 27 |
| 4.7 | Optical microscopy of the second Si/a-SiC blade after Etch 2 (SF_6 , 1 min; nominal undercut target $L \approx 12\ \mu\text{m}$). (a) Post-etch overview showing the lid-defined protected footprint and the exposed annulus outside the lid. The red box indicates the region magnified in (b). (b) Close-up of the protected-exposed transition, where a distinct contrast change indicates partial degradation of the a-SiC film in the exposed region while the protected area retains its original appearance. (c) Outer perimeter microscopy showing a released a-SiC cutting-edge overhang with length $L \approx 11\ \mu\text{m}$. Scale bars: 1 mm (a), $100\ \mu\text{m}$ (b), and $50\ \mu\text{m}$ (c). | 28 |

| | | |
|-----|--|----|
| 4.8 | Optical microscopy of released SiO ₂ cutting-edge rims at the outer perimeter of three Si/SiO ₂ blades etched with nominal undercut targets of $L \approx 60\mu\text{m}$, $100\mu\text{m}$ and $120\mu\text{m}$. (a–c) Representative edge micrographs showing a continuous released oxide overhang adjacent to the Si blade body. The measured released lengths are $L = 65\mu\text{m}$, $99\mu\text{m}$, and $120\mu\text{m}$, respectively (red dimension arrows). Scale bars: $50\mu\text{m}$ | 29 |
| 4.9 | Optical microscopy of the top-surface rim near the hub bore edge for Si/SiO ₂ blades, illustrating hub-bore protection performance of the lid. (a) Representative pre-etch bore-edge morphology (baseline, prior to all SiO ₂ etches). (b) Post-etch bore edge for a blade etched to $L \approx 60\mu\text{m}$ (representative also for $L \approx 100\mu\text{m}$), showing no obvious etch attack at the accessible top-surface rim. (c) Post-etch bore edge for the $L \approx 120\mu\text{m}$ blade, for which lid loss was reported in the etch cycle, showing clear morphological change consistent with etch exposure at the rim. Insets provide magnified views of the indicated regions. Scale bars: $25\mu\text{m}$ | 31 |
| 5.1 | Overview of the approach configuration used for the cut demonstrations, showing how optical access and controlled sample translation were achieved. (a) Top view indicating the axial camera placement relative to the blade rim, enabling visual clearance indication during the stepwise approach. (b) Front view of the blade–sample interface, illustrating the unobstructed line-of-sight to the primary contact zone at the cutting edge. (c) Isometric view of the integrated assembly (spindle and clamped blade, diagonal sample holder, and motorised linear stage), with the sample position indicated in green. | 34 |
| 5.2 | Summary image for the rigid polymer trial (Test 5, LEGO/ABS) highlighting the pre-test mark and the measured main damage-band width $w = 547\mu\text{m}$, comparable to the blade stack thickness and consistent with over-travel and substrate sidewall involvement. The full micrographs are provided in Appendix G (Fig. G.5). | 35 |
| 5.3 | Post-test optical micrographs of the controlled incremental rigid-sample trial (Test 6, LEGO/ABS) at $N \approx 17.7\text{kRPM}$. (a) Overview showing a continuous kerf of length $\approx 1496\mu\text{m}$ and central width $\approx 46\mu\text{m}$. (b,c) Higher-magnification views of the two kerf ends (boxes in (a)), showing a narrow initial kerf ($\sim 5\mu\text{m}$) that widens towards the centre; cutting debris is visible locally (e.g. in (b)). Scale bars: $100\mu\text{m}$ (a) and $20\mu\text{m}$ (b,c). | 36 |
| A.1 | M_{recon} selected by the convergence criterion of (A.7) as a function of RPM for the axial blade-radii dataset. Each point corresponds to one 7 s record plotted at its median RPM; n denotes the number of records per radius. | 45 |
| A.2 | Harmonic distortion ratio $\text{HDR}_{2\dots M_{\text{recon}}}$ (A.8) versus RPM for the axial blade-radii dataset. Each point corresponds to one 7 s record plotted at its median RPM; n denotes the number of records per radius. | 45 |
| B.1 | Operational displacement at the shaft axis ($r = 0\text{ mm}$) for baseline configurations without the blade. Solid curves denote RSR; dashed curves denote WCR_{95} (95th percentile of per-revolution peak-to-peak values within a record). Each point corresponds to one 7 s record plotted at its median RPM; n indicates the number of records per configuration. | 48 |
| B.2 | Full axial blade-radii operational displacement results. (a) RSR p–p versus rotational speed for $r = 15.5\text{ mm}$ to 17.5 mm . (b) WCR_{95} p–p versus rotational speed for the same radii. Each point corresponds to one 7 s record plotted at its median RPM; n indicates the number of records per radius. The vertical reference line indicates the minimum cutting-edge-speed requirement (15.9 kRPM , corresponding to 30 m/s at $r = 18\text{ mm}$). | 49 |
| B.3 | Fractional amplitude contribution of integer order components versus rotational speed for axial LDV measurements at multiple blade radii ($r = 16.0\text{ mm}$, 16.5 mm , 17.0 mm , 17.5 mm). For each 7 s record, the order-tracked displacement spectrum was decomposed into integer orders $k = 1 \dots M$, and the plotted quantity is the normalised contribution $ X_k / \sum_{i=1}^M X_i $ (computed over the same order range). Each marker corresponds to one record plotted at its measured median RPM; marker shape encodes radius. | 50 |
| B.4 | Per-order displacement amplitude versus rotational speed at the blade rim ($r = 17.5\text{ mm}$, axial LDV). Curves show the absolute amplitudes $ X_k $ of integer orders $k = 1 \dots M$ (with M chosen per record according to the convergence criterion described in Section A.4). Each marker corresponds to one 7 s record plotted at its measured median RPM. | 50 |

| | | |
|-----|--|----|
| C.1 | Qualitative effect of laser power on the top-edge morphology in LS-Precess mode. Keyence VHX-6000 micrographs of representative kerf segments cut at constant scan speed (200 mm s^{-1}), outlines (two), and outline pitch (0.01 mm). (a) 60 %, (b) 70 %, (c) 80 %. Red curves indicate manually traced kerf edges. Scale bars: $20 \mu\text{m}$ | 51 |
| C.2 | Schematic of blade mounting using a vacuum tweezer. The suction cup avoids contact with the outer edge but provides a compliant, single-point grip, which limits controllable alignment and axial assembly force as the bore engages the shaft seat. | 52 |
| D.1 | Schematic diagram of residual stress in a thin film of thickness d , decomposed into a uniform component σ_0 and a linear gradient component σ_1 (adapted from [53]). | 53 |
| D.2 | Curvature of a free-standing thin film of thickness d produced by a stress gradient. The sketch shows the approximate stress profile and the corresponding relaxed strain ϵ_x across the film thickness (adapted from [44]) | 54 |
| D.3 | Predicted centrifugal stiffening k_{eff}/k_0 versus rotational speed for the a-SiC design case ($w = 10 \mu\text{m}$, $d = 96 \text{ nm}$, $r_o = 18 \text{ mm}$). The equivalent deflection ratio is $h(\Omega)/h_0 = 1/(k_{\text{eff}}/k_0)$, and the relative reduction is $\Delta h/h_0 = 1 - 1/(k_{\text{eff}}/k_0)$ | 58 |
| G.1 | Representative imaging trials near the cutting edge used to assess feasibility of reducing motion blur for approach monitoring: (a) stationary rim under continuous illumination, (b) rotating rim under continuous illumination (motion blur but stable view for qualitative clearance), (c) fast-shutter trial at high ISO (rolling-shutter distortion and residual blur), and (d) strobed illumination without phase-lock (apparent drift due to phase slip and banding/artefacts under rolling-shutter video). These observations motivated the use of imaging only for qualitative clearance indication during stepwise approach. | 67 |
| G.2 | 3D-printed sample holder used for the thread cut demonstrations. (a) Isometric CAD view of the holder geometry and mounting features (camera direction indicated). (b) Front view illustrating the diagonal thread sample orientation (green) used to maintain an unobstructed axial line-of-sight to the blade-sample contact region during approach (cf. Fig. 5.1). | 68 |
| G.3 | Optical micrographs of the polyester sewing thread used as a qualitative contact indicator during the initial cut demonstrations. (a) Higher-magnification view of intact fibres prior to cutting. (b) Thread after Test 2, showing local fibre separation consistent with partial severing. (c) Thread after Test 3, showing a larger separated region consistent with further severing. Scale bars: $20 \mu\text{m}$ (a) and $100 \mu\text{m}$ (b–c). | 68 |
| G.4 | Optical micrographs (Keyence VHX-6000) of the released SiO_2 rims before and after the tensioned sewing-thread tests (Tests 2–4), used to assess edge condition and potential contact localisation. (a–d) Front-side rim on the polished wafer side (thin oxide, $\sim 4 \mu\text{m}$): (a) representative intact region with nominal protrusion length ($L \approx 120 \mu\text{m}$); (b) pre-existing local defect region prior to cutting (extent $\approx 902 \mu\text{m}$ along the rim); (c) the same region after cutting, showing an unchanged profile within the inspection resolution; (d) example of a severe local breakaway observed after the final cut attempt (origin uncertain). (e–f) Rear-side rim on the unpolished wafer side (thicker oxide, $\sim 10 \mu\text{m}$): (e) example region showing slight local wear; (f) example region showing no apparent change. Scale bars: $100 \mu\text{m}$ (a–d) and $20 \mu\text{m}$ (e–f). | 70 |
| G.5 | Post-test optical micrographs of a rigid polymer trial (white LEGO/ABS, Test 5) at $N \approx 16.9 \text{ kRPM}$. Boxed regions indicate the fields of view shown at higher magnification (a→b, b→c). (a) Overview of the modified surface region. (b) Apparent damage-band width $w = 547 \mu\text{m}$, comparable to the blade stack thickness ($\sim 539 \mu\text{m}$) and consistent with over-travel with substrate sidewall involvement (potential polymer pile-up and projection bias). The separate pre-test mark resulted from a prior touch with a clean Si blade. (c) Detail near the track end showing a local initial width of $\sim 10 \mu\text{m}$, compatible with the rear-side rim thickness but not definitive. Scale bars: $500 \mu\text{m}$ (a,b) and $20 \mu\text{m}$ (c). | 71 |
| G.6 | Representative white-light interferometry (WLI) measurement of the Test 6 kerf on ABS. The top-left panel shows the height map of a region spanning the kerf; the top-right panel shows an example line profile across the trench. Depth is reported relative to the local mean surface plane outside the kerf. In this representative scan, the profile indicates a maximum trench depth of $\Delta z \approx 14.7 \mu\text{m}$ over a lateral span of $\Delta x \approx 68.5 \mu\text{m}$. The measurement supports the depth range reported in the main text for repeated scans near the kerf centre. | 72 |

- G.7 Representative post-test microscopy of the rear-side rim after the controlled rigid-sample trial (Test 6). (a) Example region showing no apparent change within the inspection resolution. (b) Example region showing slight local degradation (increased edge roughness) over a limited sector. The coexistence of degraded and intact regions indicates spatially localised wear, consistent with intermittent contact limited to a circumferential sector. Scale bars: 50 μm 72

LIST OF TABLES

| | | |
|-----|--|----|
| 2.1 | Summary of intended bearing-seat fits, measured diameters, and implied fit conditions for the first and second manufacturing iterations. Implied-fit ranges are computed using SKF bearing tolerance bands (see notes). | 9 |
| 2.2 | Summary of the measurement configurations and operating ranges used to characterise quasi-static run-out (TIR) and operational displacement (RSR/WCR). Each operating point corresponds to one 7 s time record; results are plotted against the median RPM of that record. | 12 |
| 2.3 | Observed quasi-static total indicated run-out (TIR) at the blade outer region ($r \approx 18$ mm), measured with a 1 μ m resolution dial indicator under slow manual rotation. Values summarise the ranges observed across repeated disassembly and re-mounting (Pret ON + Clamp ON). | 13 |
| 4.1 | Summary of isotropic SF ₆ plasma etch trials performed with the lid-base fixture. Undercut lengths are nominal values based on the process calibration (approximately 12 μ m per 60 s) unless stated otherwise. | 27 |
| D.1 | Calculated initial tip deflection h_0 (in nm) as a function of cantilever length L and stress gradient $\Delta\sigma/d$, obtained from (D.5) for a-SiC with $E = 210$ GPa. | 55 |
| D.2 | Calculated dimensionless initial tip deflection h_0/L (in %) as a function of cantilever length L and stress gradient $\Delta\sigma/d$, obtained from (D.5) for a-SiC with $E = 210$ GPa. | 56 |
| D.3 | Calculated dimensionless initial tip deflection h/L (in %) as a function of cantilever length L and stress gradient $\Delta\sigma/d$. Values are obtained from (D.5) for an SiO ₂ film $E = 70$ GPa. | 56 |
| D.4 | Centrifugal stiffening factor k_{eff}/k_0 and corresponding relative tip-deflection reduction $\Delta h/h_0$ at $N = 20000$ RPM for the a-SiC design case ($w = 10$ μ m, $d = 96$ nm, $r_o = 18$ mm). | 58 |

1

INTRODUCTION

Cutting tools predate modern humans, with the oldest known cutting tools dating back approximately 3.3 million years [1], and evidence of systematic tool use persisting throughout early pre-human and human history [2]. With the development of metallurgy, purpose-made saws emerged. One of the earliest documented pull saws dates back to the Early Dynastic Period (3100–2686 BC) [3]. Much later, rotary sawing, exemplified by the circular saw, introduced continuous high-speed cutting and, after 19th-century refinements, became a widely adopted tool [4], [5], [6].

In contrast, modern micro- and nanofabrication has largely miniaturised fields and particles (beams, plasmas, and chemical reactions) rather than mechanical cutting tools. While advanced lithographic and plasma-based techniques routinely produce sub-micrometre features, deliberate miniaturisation of mechanically defined cutting implements to comparable thickness scales has remained rare. The idea of making progressively smaller tools to fabricate progressively smaller structures is conceptually compelling, yet at the nanoscale practical challenges in stiffness, handling, and control have generally favoured non-mechanical approaches [7]. This thesis revisits the macroscopic principle of rotary cutting under micro- and nanoscale constraints and investigates whether a mechanically defined edge can be realised at thin-film thickness scale.

Plasma etching, ion-beam milling, and related techniques form the backbone of contemporary nanofabrication. However, their effectiveness depends on specific physical and chemical conditions. Plasma etching requires suitable volatile reaction products and controllable ion transport [8]. Ion-beam milling depends on sufficient sputter yield and careful management of surface charging and redeposition [9]. In certain material classes, these assumptions become restrictive. In magnetically biased systems such as permanent-magnet MEMS, static magnetic fields can alter plasma transport and sheath behaviour, and more generally the fine patterning of thick permanent magnets remains difficult in practice, with no well-established MEMS-compatible route at \sim micrometre-scale thickness and lateral resolution and with existing etch-based approaches still limited by redeposition, tapering, lateral over-etch, and process-integration constraints [10], [11], [12], [13]. Chemically inert polymers may require aggressive plasma chemistries with limited selectivity, while multilayer stacks containing insulating regions can suffer from charge accumulation under ion bombardment [8], [9]. These limitations do not universally preclude beam- or plasma-based processing, but they highlight conditions under which material removal governed primarily by contact mechanics and relative hardness, rather than by chemical selectivity or sputter yield, may offer a complementary structuring route.

Mechanical micro- and nano-machining already supports demanding applications in optics, photonics, and microfluidics, but the literature reveals a persistent trade-off between minimum kerf width (the width of the trench produced by the cutting edge) and achievable trench aspect ratio (cut depth-to-width). Single-point diamond micro-tools, often shaped by focused ion beam (FIB) machining, are widely used to generate optical and fluidic microfeatures, such as diffractive optical elements and microfluidic channels. Reported tool sizes span from hundreds of micrometres down to single-digit micrometres, including spherical tips with nose radii of $7.73\ \mu\text{m}$ [14]. V-shaped tools extend this capability to grooves and rib arrays, for example producing ribs $\sim 10\ \mu\text{m}$ wide and $\sim 29\ \mu\text{m}$ deep (feature aspect ratio ~ 3) [15], [16], [17]. For photonics-type structures, rectangular-faced diamond tools have been used to machine waveguide and hologram geometries at $\sim 10\ \mu\text{m}$

tool widths [18], [19]. When deep slotting is required, stiffness becomes dominant. Specialised micro-slotting tools (e.g. 15 μm wide and 150 μm high) can reach aspect ratios on the order of 10 by using enlarged support geometries [20]. At the other extreme, Sun *et al.* demonstrated a FIB-fabricated diamond “grating” tool where individual tooth dimensions reach a base width of 459 nm, narrowing to 153 nm over a height of 594 nm, proving that sub-micrometre cutting elements are feasible in principle [21].

Across these examples, a consistent pattern emerges: as cutting elements become thinner, achievable trench depth typically decreases unless additional support geometry increases tool stiffness. Extremely thin cutting features are generally associated with limited cutting depth, whereas tools capable of deep trenching rely on thicker cross-sections or reinforced geometries to maintain structural integrity under load. Minimum kerf width and achievable trench depth are therefore strongly coupled in conventional mechanical micro-cutting approaches.

Rotating approaches face a related constraint. In micro-milling, tool diameter sets a practical floor on trench width. Representative demonstrations include $\sim 50 \mu\text{m}$ wide grooves at depths up to $\sim 180 \mu\text{m}$ (aspect ratio ~ 4) using micro-milling tools with diameters of tens of micrometres [22], [23]. The closest rotating analogue to a circular saw is the diamond dicing blade, which enables high-aspect-ratio, continuous micro-structuring including tall micro-pillar formation and deep, narrow grooves for optical waveguide geometries [24], [25]. For instance, Loiko *et al.* report diced channel waveguides $\sim 200 \mu\text{m}$ deep with 10 μm to 50 μm widths (aspect ratios 4 to 20) [25]. However, minimum kerf width in rotating blades is fundamentally constrained by blade thickness, which typically sets the lower bound on the trench width produced during cutting and is on the order of 10 μm in the surveyed literature [26]. Continuous rotating trenching with kerf widths below approximately 10 μm is therefore not evident in the surveyed literature and is typically realised via lithographic patterning and etching [27].

Taken together, these observations suggest that while mechanical approaches can achieve substantial aspect ratios and continuous trenching, extension of rotating cutting into the thin-film thickness regime would represent a distinct operating point within the broader fabrication landscape. By defining the cutting edge through thin-film deposition while relying on the bulk substrate and rotational stiffening for structural support, the wafer-based blade concept seeks to partially decouple kerf width from achievable trench depth within a rotating architecture.

To explore this possibility, this thesis proposes a wafer-based circular cutting blade in which a circumferential thin-film overhang forms an ultra-thin, continuous cutting edge around the outer perimeter. The core idea is to decouple two geometric limits that are typically linked in conventional cutting tools: the effective kerf width is defined primarily by the thin-film thickness, whereas the achievable cutting depth is governed by the film’s radial overhang length. In contrast to conventional tool fabrication routes such as focused ion beam machining, the proposed blade is developed by adapting standard semiconductor processes and selective substrate removal steps to realise the cutting geometry.

Fig. 1.1 summarises the fabrication concept. Starting from a silicon wafer carrying a thin film, the wafer is first shaped into an annular disk. In a subsequent release step, silicon is selectively removed from the outermost ring of the annulus, creating a circumferential thin-film overhang at the blade perimeter. During high-speed rotation, the rim is subjected to tangential (hoop) stress, which leads to centrifugal stiffening. Increasing rotational speed to raise the effective dynamic stiffness and suppress bending is a conventional strategy for ultra-thin dicing blades, where blades can reach rotational speeds of tens of thousands of RPM [28], [29]. In the present work, this effect is used as a design rationale and its expected trend is estimated analytically (Chapter 4). Experimentally, mechanical feasibility is evaluated via survival of the substrate and the released cutting edge at speed, and by measurements of system motion (quasi-static run-out and operational displacement).

The objective of this thesis is not to establish a process-ready alternative to lithography or plasma etching, nor to demonstrate attributable sub-micrometre kerf formation. Rather, it is to assess the feasibility of wafer-defined, thin-film rotating edges as a mechanically robust platform. Specifically, the work demonstrates (i) a repeatable fabrication route for annular blade blanks and released thin-film rims, (ii) operation of the blade at the minimum cutting-edge speed with controlled run-out, and (iii) quantified rigid kerf formation in a non-semiconductor polymer (ABS) under incremental approach. Material-specific advantages over established beam- and plasma-based methods remain open questions.

To structure this feasibility study, the thesis adopts a challenge-led framework. Each subsequent chapter




Figure 1.1: Fabrication concept for the wafer-based circular blade (schematic, not to scale). A thin film on a silicon substrate is shaped into an annular disk, after which selective substrate removal near the outer edge releases a circumferential thin-film overhang that forms a continuous cutting edge. The highlighted region indicates the substrate removal zone.

addresses one governing technical challenge: achieving stable high-speed rotation with acceptable run-out; producing smooth, continuous blade perimeters in Si/thin-film stacks; enabling selective edge release without compromising mounting geometry; and integrating the system to demonstrate first cutting engagement. The thesis concludes by synthesising the evidence across these challenges and outlining the steps required to translate the demonstrated platform towards reliable, attributable ultra-thin cutting and systematic exploration of material classes that may benefit from mechanically defined structuring.

2

CHALLENGE 1: ROTATION TO CUTTING-EDGE SPEED WITH ACCEPTABLE RUN-OUT

The first challenge addressed in this work is to develop a test setup that can reach the minimum cutting-edge speed while maintaining acceptable quasi-static run-out (TIR) and operational displacement at the blade edge. In micro- and nanoscale cutting, higher cutting-edge speeds generally improve cut quality [30], [31]. For wafer dicing in particular, increasing cutting speed reduces dicing forces and can mitigate chipping, thereby improving surface quality [29], [32]. Reported cutting-edge speeds for dicing blades, whose operating principle is most comparable to the present wafer-based blade, typically fall in the range of 30 m/s to 80 m/s [33], [34], [35], [36], [37], [38]. Therefore, this work targets 30 m/s as the minimum cutting-edge speed for initial proof-of-concept testing.

With an outer blade radius of $r_o = 18$ mm (Chapter 3), achieving 30 m/s requires an operating speed of approximately 15.9 kRPM. At 16000 RPM, the centrifugal acceleration at $r_o = 18$ mm is approximately 5100 g. A first-order estimate of centrifugal stress can be obtained from the classical solution for a thin rotating annular disk, where the maximum tangential (hoop) stress occurs at the inner radius [39]:

$$\sigma_{\theta, \max} = \frac{3 + \nu}{4} \rho \omega^2 \left(b^2 + \frac{1 - \nu}{3 + \nu} a^2 \right) \Big|_{a=r_i}, \quad (2.1)$$

where ν and ρ are Poisson's ratio and the density of the disk material, respectively; ω is angular velocity, and a and b denote the inner and outer radii, respectively. While the ideal tensile strength of single-crystal silicon is on the order of 22 GPa [40], reported practical strengths of silicon wafers can be reduced to 100 MPa to 500 MPa due to defects and processing-induced damage [41], [42], [43]. Using (2.1) with $\rho = 2329 \text{ kg m}^{-3}$, $b = r_o = 18$ mm, and treating silicon as effectively isotropic with $\nu = 0.25$ [44], yields a maximum hoop stress of approximately 2.1 MPa at 16000 RPM (evaluated near the thin-ring limit, $r_i \rightarrow r_o$). This corresponds to a conservative safety factor of roughly 47 when compared to 100 MPa. Hence, under idealised assumptions, the silicon substrate remains well below its failure point due to pure centrifugal loading.

In practice, however, fracture risk during high-speed operation is far more likely to be governed by stress concentrations at the inner edge, surface and edge flaws, and stresses introduced by clamping, imbalance, and out-of-plane vibration. These effects can locally amplify stresses far above nominal centrifugal levels and can therefore dominate failure. A key question addressed in this research is: can the silicon substrate and ultra-thin cutting edge be rotated to cutting-edge speeds while maintaining sufficiently small motion to avoid fracture and enable subsequent cutting trials (Challenge 4)?

2.1. PRACTICAL CHALLENGES IN ACHIEVING MICROMETRE-LEVEL RUN-OUT AT HIGH ROTATIONAL SPEEDS

At the target operating speeds (approximately 16 000 RPM to 20 000 RPM), small geometric errors and assembly imperfections can produce significant dynamic responses. Even when cutting forces are very low, imbalance, bearing clearance, and fit-induced misalignment can generate rotation-synchronous operational displacement and vibration that scale with rotational speed and may dominate the motion of the cutting edge. The core difficulty is therefore not generating speed per se, but doing so while maintaining a stable, repeatable rotation axis.

As a point of reference, commercial wafer-dicing spindles are commonly specified at sub-micrometre total indicated run-out (TIR), often quoted as $< 1 \mu\text{m}$ [45], [46], [47]. These specifications reflect industrial machines that rely on precision-grade spindles and bearings (e.g., air bearings and matched precision angular-contact bearings), and tight manufacturing control. By contrast, the present setup was intentionally developed as a proof-of-concept using self-manufactured structural parts and readily available components. Consequently, the objective is not to match industrial spindle performance, but to reach the target cutting-edge speed while keeping edge motion sufficiently small and repeatable for survival testing and initial cutting demonstrations.

Accordingly, this work adopts a mechanism-based acceptance criterion: the operational displacement at the blade edge should remain on the same order of magnitude as the thin-film overhang length L (typically $O(10^2 \mu\text{m})$ in this thesis). Motions much larger than L would impose alternating bending of an ultra-thin brittle edge, increasing the likelihood of fracture during free spinning or cutting. The resulting performance is therefore reported together with the realised overhang lengths, rather than judged against an industrial spindle specification.

For clarity, the dominant contributors to edge motion are grouped into three categories:

1. **Geometric run-out sources:** clearance and eccentricity in the blade bore–shaft seat interface (see Chapter 4), concentricity errors between the shaft bearing seats and blade seat, and geometric errors in the blade itself (see Chapter 4).
2. **Support compliance sources:** bearing stiffness and internal clearance, housing stiffness, and unsupported shaft overhang between the outboard bearing and the blade.
3. **Dynamic excitation sources:** mass unbalance (with forcing that scales with ω^2), motor-induced excitation, and coupling misalignment.

In later sections, edge motion is quantified using a repeatable synchronous operational displacement (RSR) and a per-revolution severity metric (WCR) (definitions in Section 2.5.2), which together separate repeatable rotation-synchronous motion from non-synchronous and time-varying components.

2.2. DESIGN APPROACH AND BASELINE CONFIGURATION

Given the limited opportunity for multiple redesign cycles and the proof-of-concept nature of this work, the mechanical setup was designed as a conservative stiffness-first baseline rather than an optimised spindle. The architecture is intentionally simple: a brushless DC motor drives a stainless steel shaft through a flexible jaw coupling; the shaft is supported by two rolling bearings housed in a rigid cylindrical housing; and the blade is mounted close to the outboard bearing to minimise bending sensitivity. A schematic overview of the assembly is provided in Fig. 2.1 (including a cross-section of the bearing housing and the blade clamping concept).

Within this architecture, the design effort focused on reducing compliance and geometric sensitivity at the blade location. First, the bearing housing was designed with a large outer diameter and short load paths to maximise stiffness at the bearing seats. Second, the blade was positioned as close as practical to the outboard bearing, reducing the effective overhang between the last support and the cutting edge. Third, stainless steel was selected for the shaft and housing as a robust baseline material that is readily machinable in the faculty workshop and provides high stiffness relative to aluminium alternatives. Stainless steel's corrosion resistance also supports potential future operation in coolant-assisted environments. Finally, the assembly was mounted to a thick base plate rigidly coupled to a granite table to reduce motion of the supporting structure.




Figure 2.1: Mechanical layout of the high-speed rotation test setup used in Challenge 1. A brushless DC motor drives a stainless steel shaft via a flexible jaw coupling; the shaft is supported by two deep-groove ball bearings in a rigid cylindrical housing mounted to an aluminium base plate. The inset section view highlights the locating/non-locating bearing concept (axial constraint by one bearing and axial float by the other) and the blade clamp arrangement, with the blade positioned close to the outboard bearing to minimise unsupported overhang.

2.2.1. BALANCING STRATEGY

Professional dynamic balancing of the complete rotating assembly was considered but not pursued. The rotating stack is frequently assembled and disassembled during experimental work, and small changes in assembly or small damage to parts can alter the mass distribution. Therefore, high-grade balancing of a single configuration would not necessarily carry over between tests. To illustrate the practical implication, the standard balance-grade relation $e_{\text{per}} = G/\omega$ implies that at 20 000 RPM ($\omega \approx 2094 \text{ rad s}^{-1}$) a balance quality grade of $G = 2.5 \text{ mm s}^{-1}$ (e.g., high-performance equipment, gas turbines, high-speed armatures) corresponds to a permissible residual eccentricity of only $\sim 1.2 \mu\text{m}$, while $G = 0.4 \text{ mm s}^{-1}$ (e.g., precision machinery, high-speed dicing spindles, and gyroscopes) corresponds to $\sim 0.2 \mu\text{m}$. For the $\sim 145 \text{ g}$ rotating mass of the present shaft–blade stack, this corresponds to residual unbalance on the order of 0.2 g mm and 0.03 g mm , i.e. a demanding requirement for a repeatedly reassembled research setup. The approach taken was therefore to prioritise geometric quality (fits, concentricity, and repeatable assembly) and to quantify the resulting motion experimentally.

2.3. BEARING SELECTION AND AXIAL CONSTRAINT CONCEPT

A central design choice is the bearing concept, because bearing stiffness, internal clearance, and fit quality strongly influence run-out at high speed. The present application operates with a rotating inner ring, stationary outer ring, and an approximately constant load direction (gravity-dominated, with negligible cutting force). Under these conditions, the load is *rotating* with respect to the inner ring and *stationary* with respect to the outer ring, which is the classical case where the inner ring is typically given a tight fit to prevent creeping [48, Table 1]. At the same time, the applied load is relatively light (dominated by the shaft weight), and SKF notes that ensuring a minimum load becomes increasingly important at high speed, particularly when operating above $\sim 50\%$ of the limiting speed [48, p. 106]. For this reason, the bearing selection was biased towards a deep-groove ball bearing with a limiting speed comfortably above the operating window.

Two bearing concepts were considered: (i) a cost-effective baseline using high-quality deep-groove ball bear-

ings, and (ii) a high-performance design using a high-precision preloaded matched set of angular-contact bearings. While high-precision angular-contact bearings offer high stiffness, ultra-low axial and radial run-out, and controllable preload, they are substantially more costly. Accordingly, the baseline setup adopted deep-groove ball bearings (SKF 6000 series). For the selected bearing size, the SKF catalogue lists limiting speeds on the order of 40 000 RPM for the open bearing (and notes that single-shield variants use the open-bearing limiting speed) [48]. This places the operating window (16 000 to 20 000 RPM) at roughly 40–50% of the limiting speed, which supports stable high-speed operation provided that fits and lubrication are adequate.

The bearings were implemented using a locating/non-locating constraint strategy: one bearing axially locates the shaft, while the second allows controlled axial freedom to accommodate thermal growth and avoid over-constraint. This axial constraint concept is illustrated in the inset of Fig. 2.1. This choice also drives the fit strategy. SKF recommends interference or transition fits in precision or high-speed applications to minimise deflections and vibration [48, p. 143], motivating an interference/transition fit on the shaft for both inner rings (targeted here via a j6 shaft seat). For the housing, the locating bearing was targeted with a transition fit (J6) to provide radial support while allowing assembly. For the non-locating bearing, axial sliding must be possible. SKF therefore notes that when a non-locating bearing needs to move axially on its seat, the ring subjected to stationary load should have a loose fit [48, p. 143]. In the present configuration this implies a loose fit for the outer ring in the housing (targeted via H7). Practically, this is a controlled compromise where too tight a fit risks preventing axial float and introducing unintended preload at speed, whereas too loose a fit can increase radial freedom and contribute to radial run-out. The as-built fits and their implications are evaluated below.

2.4. MANUFACTURING OF BEARING SEATS AND FIT VERIFICATION

Achieving micrometre-level bearing fits and adequate coaxiality using conventional workshop lathes (rather than CNC or ultra-precision equipment) was a dominant practical challenge in developing the rotation test setup. High-speed run-out is strongly influenced by the geometric quality of the bearing seats (diameter, roundness, surface finish) and by the coaxiality between the bearing seats and the blade seat. Workshop guidance indicated that the positioning capability and readout resolution are $O(10\mu\text{m})$, whereas the intended fit conditions require control at the single- to few-micrometre level. Consequently, manufacture of the shaft bearing seats and housing bores became iterative, and early iterations did not meet the intended fits.

Direct metrology for coaxiality between seats, roundness, and quantitative surface finish was not available within the workshop. Therefore, beyond diameter checks, these properties were evaluated indirectly by treating the assembled rotor operational displacement as an integrated performance metric that captures the combined effects of geometry, fits, and support compliance.

Table 2.1 shows why the first manufacturing iteration was not suitable for high-speed testing. On Shaft 1, the locating bearing seat was undersized, resulting in a clearance-dominated condition. At high speed, this increases the risk of micro-slip (creep) between ring and seat, which can lead to fretting debris, local heating, and increasing run-out. In addition, the non-locating housing bore of Bearing housing 1 was substantially oversized relative to H7, increasing the likelihood that the outer ring could settle with a small radial offset and contribute directly to blade run-out.

The second iteration achieved the intended fit philosophy more consistently. On Shaft 2, both bearing seats on the shaft fall within the intended interference regime, supporting secure inner-ring mounting. Bearing housing 2 provides a loose but reasonable H7-type condition for the non-locating outer ring, enabling axial float. The locating bore in Bearing housing 2 lies near the loose side of the intended transition fit. In principle this can allow micrometric ring motion under dynamic excitation and may contribute to run-out, but given the project time constraints it was accepted for the initial high-speed measurement test.

Assembly outcome (qualitative check). The achieved fits were consistent with the observed assembly behaviour. Bearings could be seated using thermal assistance (heated bearing/housing and cooled shaft), requiring at most light mechanical persuasion. During insertion, ring-support measures were used to avoid transmitting assembly loads through the rolling elements (i.e. force was applied through the intended ring rather than across the bearing races).

Table 2.1: Summary of intended bearing-seat fits, measured diameters, and implied fit conditions for the first and second manufacturing iterations. Implied-fit ranges are computed using SKF bearing tolerance bands (see notes).

| Part | Feature | Intended (mm) | Measured (mm) | Deviation (mm) | Implied fit (mm) |
|-------------------|---------------------------------|---|---------------|----------------|---|
| Shaft 1 | 1st bearing seat (non-locating) | 10.000 ^{+0.007} _{-0.002} (j6) | 9.998 | -0.002 | -0.006 to 0.002 (transition) |
| Shaft 1 | 2nd bearing seat (locating) | 10.000 ^{+0.007} _{-0.002} (j6) | 9.993 | -0.007 | -0.001 to 0.007 (transition- loose) |
| Bearing housing 1 | 1st bearing seat (non-locating) | 26.000 ^{+0.021} _{-0.000} (H7) | 26.040 | 0.040 | 0.040 to 0.049 (too loose) |
| Bearing housing 1 | 2nd bearing seat (locating) | 26.000 ^{+0.008} _{-0.005} (J6) | 25.989 | -0.011 | -0.011 to -0.002 (interference) |
| Shaft 2 | 1st bearing seat (non-locating) | 10.000 ^{+0.007} _{-0.002} (j6) | 10.004 | 0.004 | -0.012 to -0.004 (interference) |
| Shaft 2 | 2nd bearing seat (locating) | 10.000 ^{+0.007} _{-0.002} (j6) | 10.002 | 0.002 | -0.010 to -0.002 (interference) |
| Bearing housing 2 | 1st bearing seat (non-locating) | 26.000 ^{+0.021} _{-0.000} (H7) | 26.005 | 0.005 | 0.005 to 0.014 (loose) |
| Bearing housing 2 | 2nd bearing seat (locating) | 26.000 ^{+0.008} _{-0.005} (J6) | 26.003 | 0.003 | 0.003 to 0.012 (risk of looseness) |

Notes: Implied-fit ranges are computed using SKF bearing tolerance bands: *bearing bore* for 10 mm: max = 0 μm, min = -8 μm; *bearing outer diameter* for 26 mm: max = 0 μm, min = -9 μm. Positive values denote clearance; negative values denote interference.

Lubrication commissioning. The bearings were supplied without lubrication and were lubricated prior to operation. Before lubrication, hand-spinning produced clearly audible noise and perceptible vibration. After applying a light lubricant, both effects reduced substantially. Uncalibrated phone microphone indicator sound-level checks (microphone at ~15 mm) indicated peak levels decreasing from ~60 dB to the ambient level ~40 dB). This is treated as a qualitative commissioning check that guided the lubrication choice rather than as a quantitative assessment of bearing condition.

2.5. VERIFICATION OF SPEED AND DISPLACEMENT MEASUREMENTS

This section describes the measurement chain used to obtain rotational speed, quasi-static run-out (TIR), and operational displacement, and defines the key performance indices (KPIs) used to report setup performance in Section 2.6.

2.5.1. TACHOMETRY: SPEED MEASUREMENT VERIFICATION

Because the spindle is driven by a brushless motor without an integrated encoder, shaft speed was measured independently using a low-resolution tachometer signal. In addition to providing an RPM readout, the tachometer signal is used as a synchronous reference for rotation-synchronous analysis of the operational displacement measurements reported later in this chapter. Fig. 2.2 summarises the integrated sensing configuration.

SENSOR SELECTION AND PULSE FREQUENCY

In high-end spindles, speed is typically obtained from optical or magnetic encoders with many counts per revolution. Here only speed and a phase reference were required; therefore a low-resolution approach with $N_{\text{ppr}} = 1-2$ pulses per revolution was sufficient. The pulse frequency is

$$f = \frac{N_{\text{ppr}} \text{RPM}}{60}, \quad (2.2)$$

which at RPM = 20000 yields $f \approx 333$ Hz for $N_{\text{ppr}} = 1$ and $f \approx 667$ Hz for $N_{\text{ppr}} = 2$, i.e. well below the bandwidth of common optical-reflective and Hall-switch sensors. Consequently, sensor selection was governed primarily by integration constraints (added rotating mass, adjustability during commissioning, and mounting risk). Practical non-contact options considered were optical-reflective sensing, optical-transmissive sensing (photointerrupter), and magnetic sensing (Hall-effect switch). Given the space envelope and the preference to avoid rigid attachments on the rotor, optical-transmissive sensing was discarded. Optical-reflective

sensing was selected because it allows rapid iteration of target geometry without permanent shaft modification, whereas magnets require reliable retention under centrifugal loading and offer limited adjustability once bonded.

2

ORDER-OF-MAGNITUDE CHECK: ADDED MASS AND PLACEMENT SENSITIVITY OF TACHOMETER TARGETS

Two approximately opposing windows were used to reduce added eccentricity. In practice their angular separation was positioned by eye and may deviate from 180° . If two equal target masses m at radius r are separated by $\pi + \varepsilon$ (placement error ε), the residual unbalance magnitude $U(\varepsilon)$ and the synchronous ($1\times$) force amplitude $F(\varepsilon)$ are

$$U(\varepsilon) = \left| mr + mre^{i(\pi+\varepsilon)} \right| = 2mr \left| \sin\left(\frac{\varepsilon}{2}\right) \right|, \quad F(\varepsilon) = U(\varepsilon) \omega^2, \quad \omega = 2\pi \frac{\text{RPM}}{60}. \quad (2.3)$$

This scaling shows that, for fixed r , ε , and ω , the synchronous unbalance force increases linearly with the added mass m , motivating lightweight and approximately symmetric target features at high speed.

As an illustrative example at $\text{RPM} = 20000$ with $r = 4 \text{ mm}$ and $\varepsilon = 5^\circ$, an estimated tape-segment mass of $m \approx 3 \times 10^{-2} \text{ g}$ gives $F \approx 4 \times 10^{-2} \text{ N}$, whereas a micro magnet of $m \approx 7 \times 10^{-3} \text{ g}$ would reduce this force to $F \approx 1 \times 10^{-2} \text{ N}$. Although magnets can reduce the synchronous force amplitude, they require reliable retention and offer limited adjustability once bonded to the shaft, which motivated the optical-reflective solution used here. Non-ideal 180° spacing can also introduce phase non-uniformity when deriving a $1\times$ reference from a $2\times$ tachometer signal.

SENSOR IMPLEMENTATION AND OSCILLOSCOPE CONFIGURATION

An optical-reflective TCRT5000 module was configured for $N_{\text{ppr}} = 2$ by blacking out the shaft with black tape while leaving two narrow exposed windows of bare stainless steel. For order tracking, a $1\times$ reference was obtained by triggering on every second rising edge. The sensor position is shown in Fig. 2.2 (inset). Initial trials using small IR-reflective tape patches produced a less robust trigger waveform. Therefore, the blacked-out shaft approach was adopted, and a black shroud was added around the sensor to reduce sensitivity to ambient reflections.

The tachometer output was recorded on oscilloscope Channel 1. For accurate edge timing, CH1 was configured with DC coupling. As shown in Fig. 2.3(a), AC coupling introduces baseline droop and sloped plateaus, so a fixed trigger threshold can shift in time relative to the true edge. DC coupling (Fig. 2.3(b)) preserves stable logic levels, enabling repeatable threshold triggering and a stable phase reference for synchronous averaging.


During commissioning, the oscilloscope-derived RPM was cross-checked against a handheld laser tachometer pointed at a single reflective patch on the shaft. Agreement was observed within the displayed readout resolution, providing confidence in the RPM estimate. However, no systematic calibration was recorded.

2.5.2. DISPLACEMENT MEASUREMENT VERIFICATION

Throughout this chapter, *run-out* refers to quasi-static total indicated run-out (TIR) obtained by slowly rotating the spindle by hand and recording the dial-indicator max–min over one revolution. *Operational displacement* refers to in-operation motion reconstructed from LDV velocity measurements during rotation. Two KPIs are reported from the LDV data: (i) the repeatable synchronous run-out (RSR), defined as the peak-to-peak amplitude of the rotation-synchronous reconstructed displacement waveform, and (ii) the worst-case run-out (WCR), defined as a per-revolution peak-to-peak measure summarised by the 95th percentile over all captured revolutions.

Quasi-static TIR was measured with a dial indicator (resolution $1 \mu\text{m}$) at the blade outer region using a test blade without a cutting edge. For axial TIR, the probe was positioned on the blade face as close as practical to the outer radius ($r \approx 17 \text{ mm}$ to 18 mm) and aligned normal to the surface to minimise cosine error. For radial TIR, the probe contacted the blade sidewall near mid-thickness. The shaft was rotated slowly by hand and the max–min reading over a full revolution was recorded. Measurements were repeated after disassembly and re-mounting to characterise assembly sensitivity.

Operational displacement was obtained from the measured velocity in the order domain (harmonic conversion), avoiding time-domain integration drift (Appendix A). In the axial configuration (Fig. 2.2), the LDV beam was aligned with the spindle axis. A single-axis translation stage allowed repeatable radial repositioning of the LDV spot to select the measurement radius r (e.g. $r = 0$ on the shaft centre and $r = 15.5 \text{ mm}$ to 17.5 mm on the



2

Figure 2.2: Integrated sensing configuration for tachometry and LDV measurements. An optical-reflective tachometer sensor is mounted above the 8 mm-diameter shaft (inset) to provide a pulse train used for RPM estimation and synchronous triggering (CH1). A Laser Doppler Vibrometer (LDV) measures *axial* vibration velocity and operational displacement is reconstructed from the velocity signal in post-processing. A single-axis translation stage repositions the LDV spot radially to select the measurement radius r (e.g. $r = 0$ on the shaft centre and $r = 15.5$ mm to 17.5 mm on the blade). The analogue decoder output is recorded on the oscilloscope (CH2).

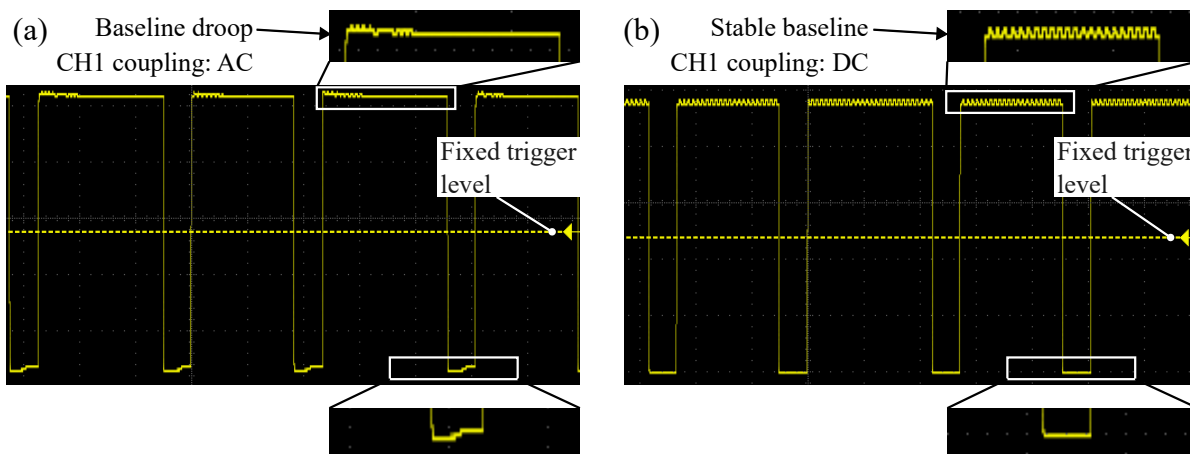


Figure 2.3: Commissioning of the tachometer trigger signal on oscilloscope CH1. (a) With AC coupling enabled, the square-wave pulse train exhibits baseline droop and sloped plateaus (insets), consistent with attenuation of low-frequency content by the AC-coupling network. A fixed trigger threshold can therefore shift in time relative to the true edge. (b) With DC coupling, the waveform preserves stable logic levels and a constant baseline (insets), enabling reliable threshold triggering. In subsequent measurements, $N_{ppr} = 2$ pulses per revolution were generated and a $1\times$ reference was obtained by triggering on every second rising edge.

blade). For radial operational displacement, the spindle base assembly was rotated by 90° such that the LDV beam impinged the blade sidewall. The spot was positioned at mid-thickness and return quality was verified using the sensor signal indicator.

2.6. SETUP PERFORMANCE: QUASI-STATIC RUN-OUT AND OPERATIONAL DISPLACEMENT VERSUS SPEED

This section quantifies setup performance in terms of (i) quasi-static run-out, expressed as total indicated run-out (TIR) measured with a dial indicator, and (ii) operational displacement measured with the LDV. Operational displacement is reported using the repeatable synchronous run-out (RSR) and worst-case run-out (WCR) metrics defined in Section 2.5.2. The primary interest is the blade rim, where the cutting edge is lo-

cated. Results are therefore reported at radii near $r \approx 18$ mm and compared against the targeted overhang lengths in the acceptance check (Section 2.6.5).

2.6.1. TEST MATRIX AND MEASUREMENT POINTS

Each data point in Section 2.6 corresponds to one recorded time trace of approximately 7 s. Because the achieved speed depends on the ESC setting and can drift slightly during a record, results are plotted against the measured median rotational speed of each record. The test matrix (Table 2.2) comprised three baseline configurations without the blade to assess sensitivity to the axial bias assembly procedure and clamp attachments at the shaft axis ($r = 0$ mm), followed by blade-mounted tests at multiple radii near the rim in the axial LDV configuration, and rim tests in the radial LDV configuration. For conciseness in the main text, the rim location is reported as $r \approx 18$ mm (axial: $r = 17.5$ mm; radial: $r = 18$ mm).

2.6.2. EFFECT OF AN ASSEMBLY-INDUCED AXIAL BIAS LOAD

No controlled bearing-preload mechanism (e.g. wave spring or calibrated spacer) was implemented in this setup. However, SKF notes that preload may be suitable when a high degree of stiffness or positional control is required [48, p. 186]. To assess whether an axial bias load could reduce operational displacement in the present assembly, an exploratory commissioning check was performed at the shaft axis ($r = 0$) using three configurations: (i) baseline with no axial bias load (Pret OFF, Clamp OFF, Blade OFF), (ii) assembly-induced axial bias load applied by pulling axially on the shaft while tightening the motor coupling (Pret ON, Clamp OFF, Blade OFF), and (iii) the same axial bias procedure with the shaft clamp installed (Pret ON, Clamp ON, Blade OFF).

The applied force magnitude was unmeasured and therefore this procedure should not be interpreted as a calibrated bearing preload. Nevertheless, the procedure was repeatable in direction (consistent pull direction) and showed a clear reduction in hub ($r = 0$ mm) RSR relative to the non-pretensioned baseline: in the overlapping mid-RPM range, RSR decreased by approximately 35–40%, and at the highest tested RPM by approximately 20–25% (Appendix B, Fig. B.1). No obvious adverse effects (e.g. audible noise increase) were observed during these short measurements, although bearing temperature was not monitored. Based on this empirical reduction in rotation-synchronous motion, the axial bias assembly procedure was used in all subsequent rotational and cutting tests reported in this thesis.

2.6.3. QUASI-STATIC RUN-OUT AND ASSEMBLY SENSITIVITY (TIR)

Quasi-static TIR was measured using the dial-indicator procedure described in Section 2.5.2. Measurements were repeated after disassembly and re-mounting to quantify assembly sensitivity; unless stated otherwise, the results below correspond to the mounting approach used during the operational LDV tests (Pret ON + Clamp ON).

Table 2.3 summarises the observed TIR ranges. In the axial direction, very low values (below $5 \mu\text{m}$) were achievable, but the outcome was highly sensitive to re-mounting. After an apparently similar re-assembly, axial TIR could increase to $\mathcal{O}(10 \mu\text{m})$ – $\mathcal{O}(60 \mu\text{m})$. Attempts to systematically relate axial TIR to the angular mounting orientation of the clamp did not reveal a consistent trend. Instead, a varying visible gap between the blade and the clamp face suggested that small, non-repeatable changes in seating and tilt dominated the axial run-out state.

Table 2.2: Summary of the measurement configurations and operating ranges used to characterise quasi-static run-out (TIR) and operational displacement (RSR/WCR). Each operating point corresponds to one 7 s time record; results are plotted against the median RPM of that record.

| Purpose | Pretension | Clamp | LDV direction / radius r | N | RPM range |
|-----------------------|------------|-------|--------------------------------------|-----|------------------|
| Baseline (axis) | OFF | OFF | Axial, $r = 0$ mm | 3 | 7.7 to 13.4 kRPM |
| Baseline (axis) | ON | OFF | Axial, $r = 0$ mm | 4 | 6.0 to 14.4 kRPM |
| Baseline (axis) | ON | ON | Axial, $r = 0$ mm | 5 | 7.1 to 13.4 kRPM |
| Blade mapping (axial) | ON | ON | Axial, $r = 15.5$ mm to 17.5 mm | 29 | 6.0 to 19.7 kRPM |
| Blade rim (radial) | ON | ON | Radial, $r \approx 18$ mm (sidewall) | 8 | 6.0 to 20.8 kRPM |

Note: “Pretension” denotes the repeatable assembly-induced axial bias load applied during tightening of the motor coupling, as defined in Section 2.6.2.

Table 2.3: Observed quasi-static total indicated run-out (TIR) at the blade outer region ($r \approx 18$ mm), measured with a $1 \mu\text{m}$ resolution dial indicator under slow manual rotation. Values summarise the ranges observed across repeated disassembly and re-mounting (Pret ON + Clamp ON).

| Direction | Min | Typical | Max |
|-------------------------------|------------------|--------------------------------------|-----------------------|
| Axial TIR, $r \approx 18$ mm | $<5 \mu\text{m}$ | $\sim 20 \mu\text{m}$ | $\sim 60 \mu\text{m}$ |
| Radial TIR, $r \approx 18$ mm | $30 \mu\text{m}$ | $30 \mu\text{m}$ to $60 \mu\text{m}$ | $60 \mu\text{m}$ |

Radial TIR values typically lay in the range $30 \mu\text{m}$ to $60 \mu\text{m}$ across re-mounting. Because this is a contact measurement on a thin sidewall, the radial TIR is treated as an upper-bound indicator of quasi-static geometric eccentricity. Local waviness/roughness and probe contact conditions can inflate the peak-to-peak reading. Importantly, quasi-static TIR should not be interpreted as equivalent to the operational displacement metrics reported later. The LDV-based RSR/WCR quantify operational displacement at a local optical spot and explicitly separate the repeatable rotation-synchronous component (RSR) from per-revolution worst-case behaviour (WCR).

Potential bias sources (TIR).

- **Local surface waviness/roughness (radial):** contact on the thin sidewall can inflate peak-to-peak readings depending on the local track.
- **Probe alignment and cross-sensitivity:** small misalignment can introduce cosine error and cross-axis coupling.
- **Seating and tilt:** small particles, burrs, or uneven contact at the clamping interface can change the tilt state between shaft and blade.
- **Contact conditions:** probe preload and contact point height (relative to the $525 \mu\text{m}$ thickness) are difficult to reproduce exactly.

2.6.4. OPERATIONAL DISPLACEMENT VERSUS ROTATIONAL SPEED (RSR AND WCR)

Operational displacement during rotation was quantified at the blade using the LDV-derived metrics defined in Section 2.5.2. Fig. 2.4 shows the repeatable synchronous run-out (RSR) and the worst-case run-out metric WCR_{95} (95th percentile of per-revolution peak-to-peak values within a record), plotted against the measured median RPM of each 7 s record.

In the axial configuration (Fig. 2.4(a)), RSR at the rim location ($r \approx 18$ mm) remained in the tens of micrometres across the required operating regime, increasing from $\approx 21 \mu\text{m}$ at 15.96 kRPM to $\approx 33 \mu\text{m}$ at 19.7 kRPM. A representative inner radius is included to indicate that operational displacement generally increases towards the rim. Full multi-radius results are provided in Fig. B.2 (Appendix B). Both RSR and WCR_{95} show a clear upturn beyond approximately 18 kRPM, which is treated as a practical indicator that the setup becomes increasingly sensitive at the cutting-edge-speed regime. This threshold also coincided with the first vibration that became clearly perceptible by hand on the granite base, supporting the interpretation that the system becomes more excitation-sensitive beyond 18 kRPM.

In the radial configuration (Fig. 2.4(b)), rim displacement showed comparable magnitudes, with RSR increasing from $\approx 16 \mu\text{m}$ at 16.7 kRPM to $\approx 21 \mu\text{m}$ at 20.8 kRPM. Across both directions, WCR_{95} typically exceeded RSR by only a few micrometres at a given speed, indicating that the operational motion is predominantly rotation-synchronous and therefore relatively repeatable on a per-revolution basis. The separation between WCR_{95} and RSR becomes more pronounced at the highest speeds, consistent with the upturn beyond 18 kRPM.

2.6.5. ACCEPTANCE VERSUS OVERHANG LENGTH

The cutting-edge overhang lengths targeted in this work are $L = 60 \mu\text{m}$, $100 \mu\text{m}$, and $120 \mu\text{m}$ (see Chapter 4). To relate Challenge 1 performance directly to these values, the measured operational displacement amplitudes at the blade rim ($r \approx 18$ mm) are normalised as the non-dimensional ratios RSR/L and WCR_{95}/L . This normalisation is conservative because RSR and WCR_{95} are peak-to-peak values (rather than one-sided amplitudes).

At the minimum cutting-edge-speed requirement (15.9 kRPM), the axial rim displacement was $\text{RSR} \approx 21 \mu\text{m}$

and $WCR_{95} = 23.6 \mu\text{m}$, corresponding to $RSR/L = \{0.35, 0.21, 0.18\}$ and $WCR_{95}/L = \{0.39, 0.24, 0.20\}$ for $L = \{60 \mu\text{m}, 100 \mu\text{m}, 120 \mu\text{m}\}$. At the highest tested axial rim speed (19.7 kRPM), the displacement increased to $RSR \approx 33 \mu\text{m}$ and $WCR_{95} = 37.2 \mu\text{m}$, giving $RSR/L = \{0.55, 0.33, 0.28\}$ and $WCR_{95}/L = \{0.62, 0.37, 0.31\}$.

In the radial rim configuration, the minimum-speed point (16.7 kRPM) yielded $RSR \approx 16 \mu\text{m}$ and $WCR_{95} = 19.5 \mu\text{m}$, resulting in $RSR/L = \{0.27, 0.16, 0.13\}$ and $WCR_{95}/L = \{0.32, 0.19, 0.16\}$. At the highest tested radial rim speed (20.8 kRPM), $RSR \approx 21 \mu\text{m}$ and $WCR_{95} = 25.4 \mu\text{m}$, giving $RSR/L = \{0.35, 0.21, 0.18\}$ and $WCR_{95}/L = \{0.42, 0.25, 0.21\}$. In all cases, both the typical synchronous motion (RSR) and the conservative per-revolution severity measure (WCR_{95}) remained below the overhang length ($RSR/L < 1$ and $WCR_{95}/L < 1$). Therefore, for the overhang range targeted in this thesis ($L = 60 \mu\text{m}$ to $120 \mu\text{m}$), the measured operational displacement at the blade rim is sufficiently small to proceed with survival testing and initial cutting demonstrations at and above the minimum cutting-edge speed.

2.7. DISCUSSION AND IMPLICATIONS

Challenge 1 establishes a practical operating envelope for the proof-of-concept setup: cutting-edge speeds of $\geq 30 \text{ m/s}$ were achieved while maintaining rim displacement in the tens of micrometres. The observed upturn in both RSR and WCR_{95} beyond $\sim 18 \text{ kRPM}$ indicates increasing sensitivity at the upper end of the tested speed range, and quasi-static TIR was found to be strongly assembly dependent.

Order-domain decomposition of the LDV displacement (amplitude fractions and per-order amplitudes) shows that the operational motion is concentrated in the lowest integer orders (Appendix B, Fig. B.3, and Fig. B.4). Across radii and operating points, the low-to-mid-speed response frequently exhibits a pronounced $2\times$ component, whereas towards the upper end of the operating window the $1\times$ component grows strongly and becomes comparable to, or larger than, $2\times$. This shift is consistent with two concurrent contributors. First, a repeatable twice-per-revolution modulation can arise from geometric or stiffness asymmetry in the clamp-shaft-bearing stack (e.g. slight tilt/misalignment or anisotropic support). Second, a speed-amplified synchronous contribution compatible with residual unbalance and clearance-dependent rotor centring can become more critical near the observed upturn above $\sim 18 \text{ kRPM}$, as its forcing scales strongly with ω^2 . This high-speed sensitivity also coincided with the first vibration that was clearly perceptible by hand on the granite base.

Motor-drive excitation cannot be ruled out, because the spindle was operated with a low-cost sensorless ESC and commutation-related torque ripple can occur at higher mechanical orders than those shown here (for six-step commutation, ~ 12 torque steps occur per mechanical revolution for a 4-pole motor). However, the present reconstruction was intentionally truncated at the smallest order range that yields a stable peak-to-peak displacement estimate, and in this dataset M_{recon} typically lay in the single-digit range (Appendix A, Fig. A.1). Therefore, while higher-order ripple may still modulate the motion, the measured rim displacement reported via RSR/WCR is likely primarily set by low-order geometric and rotor-support effects. A planned A/B comparison using a field-oriented controller was attempted but could not be completed due to hardware failure. Taken together, the most direct routes to further reduce rim displacement in future iterations are improved geometric repeatability of the blade clamping stack, implementation of a controlled preload/axial-location strategy, and dedicated characterisation or upgrading of the motor drive to quantify and reduce torque ripple at high speed.

Finally, this chapter answers the *substrate-level* part of the opening feasibility question. Stable operation of the silicon substrate was demonstrated up to at least 19.7 kRPM without macroscopic fracture during the LDV measurement campaigns. This supports the conclusion that, in the developed clamping and bearing configuration, the silicon substrate can be rotated to cutting-edge-speed conditions without catastrophic failure. The remaining uncertainty concerns the survival and engagement behaviour of the released ultra-thin edge, which is addressed in subsequent challenges.

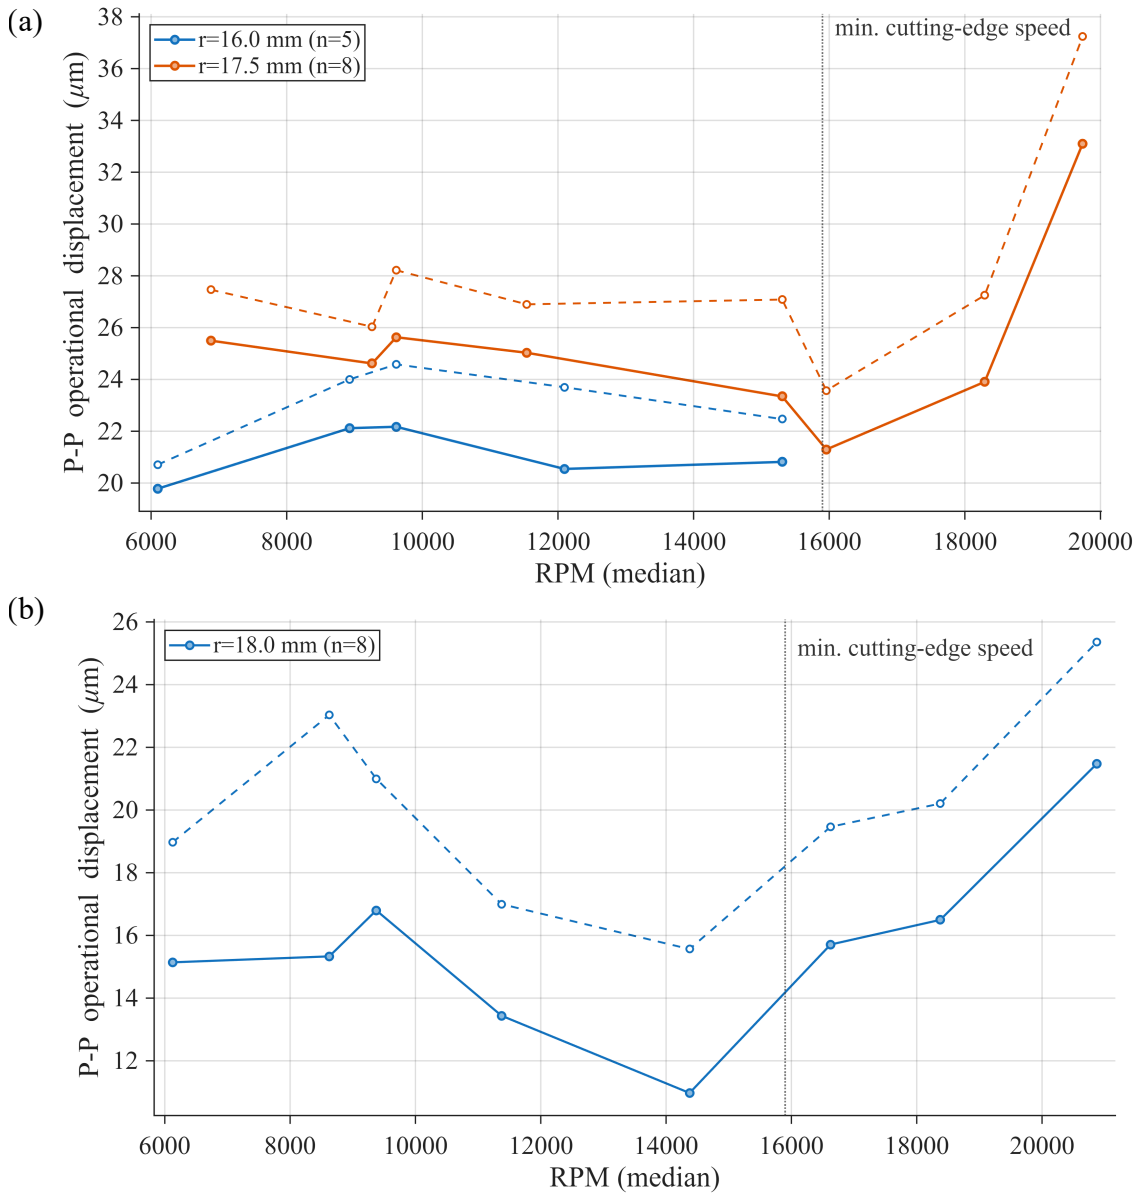


Figure 2.4: Operational displacement versus rotational speed at blade radii, showing the repeatable rotation-synchronous component (RSR) and per-revolution severity metric (WCR₉₅). (a) Axial LDV configuration at $r = 16.0\text{ mm}$ and $r = 17.5\text{ mm}$. (b) Radial LDV configuration at the rim ($r \approx 18\text{ mm}$). Solid curves denote RSR; dashed curves denote WCR₉₅. Each marker corresponds to one 7 s time record and is plotted at its measured median RPM; n in the legend indicates the number of records in each series.

3

CHALLENGE 2: FEMTOSECOND-LASER FABRICATION OF ANNULAR BLADE BLANKS

The second challenge addressed in this work is to fabricate annular blade blanks from 525 μm -thick silicon wafer stacks (Si/a-SiC and Si/SiO₂) with sufficient geometric integrity to serve as inputs to the subsequent underetching (Challenge 3) and high-speed rotation experiments (Challenge 1, 4). In this chapter, a blade blank is considered *usable* if it (i) is fully cut through without macroscopic backside exit defects (chipping or local bulge formation), (ii) exhibits an edge condition that still permits formation of a continuous ultra-thin overhang after underetching, and (iii) can be assembled repeatably on the shaft with a sufficiently small bore clearance to limit the potential for eccentric mounting.

3.1. CONSTRAINT-DRIVEN BLADE OUTLINE SELECTION

The available wafers are 100 mm in diameter, so the blade outer radius is set primarily by manufacturing yield and outline accuracy. To maximise samples for the etching trials, the design target was four blades per wafer. At $r_o \approx 20$ mm the clearance to the wafer edge (and neighbouring blades) becomes small, so modest centring errors during manual alignment in the LS-Lab software can render a blade unusable.

A second constraint is the scanner field size, which in the current LASEA LS-Lab configuration accommodates circular profiles up to $r_o = 18$ mm within a single field. Choosing $r_o = 18$ mm therefore avoids scan-field stitching, which would add positional uncertainty at the stitch boundary and reduce outline accuracy. For these reasons $r_o = 18$ mm was selected. The trade-off is a modest increase in spindle speed to reach $v_c = 30$ m/s: $N \approx 15.9$ kRPM at 18 mm versus 14.3 kRPM at 20 mm ($\sim 10\%$).

3.2. LS-PRECESS ENABLES THROUGH-CUTS IN 525 μm SILICON

The annular blanks were produced by femtosecond-laser micromachining in the LASEA LS-Lab system. Initial tests were performed to establish whether full-depth cutting through 525 μm Si is feasible within practical machining times using the standard ("normal") cutting mode, or whether the LS-Precess module is required. In LS-Precess mode, a controlled beam motion is imposed around the optical axis to improve deep-cutting performance [49]. In the present work, this module constitutes the key enabling intervention: without a reliable through-cut, subsequent optimisation of edge condition and bore-fit cannot be validated.

Fig. 3.1 compares parameter sweeps in normal mode and LS-Precess mode for identical circular test patterns. In normal mode, no through-cuts were obtained within the explored power–repetition window (Fig. 3.1(a)), and even an extreme setting of 80% power with 60000 repetitions (otherwise identical scan settings) did not produce a reliable through-cut. In contrast, LS-Precess mode produced through-cuts across all tested points in the explored parameter window (Fig. 3.1(b)). Importantly, the top-surface appearance alone is not a reliable indicator of cut completion. Both modes can exhibit similar near-edge surface features, while only LS-Precess yielded full-depth separation under the tested conditions.

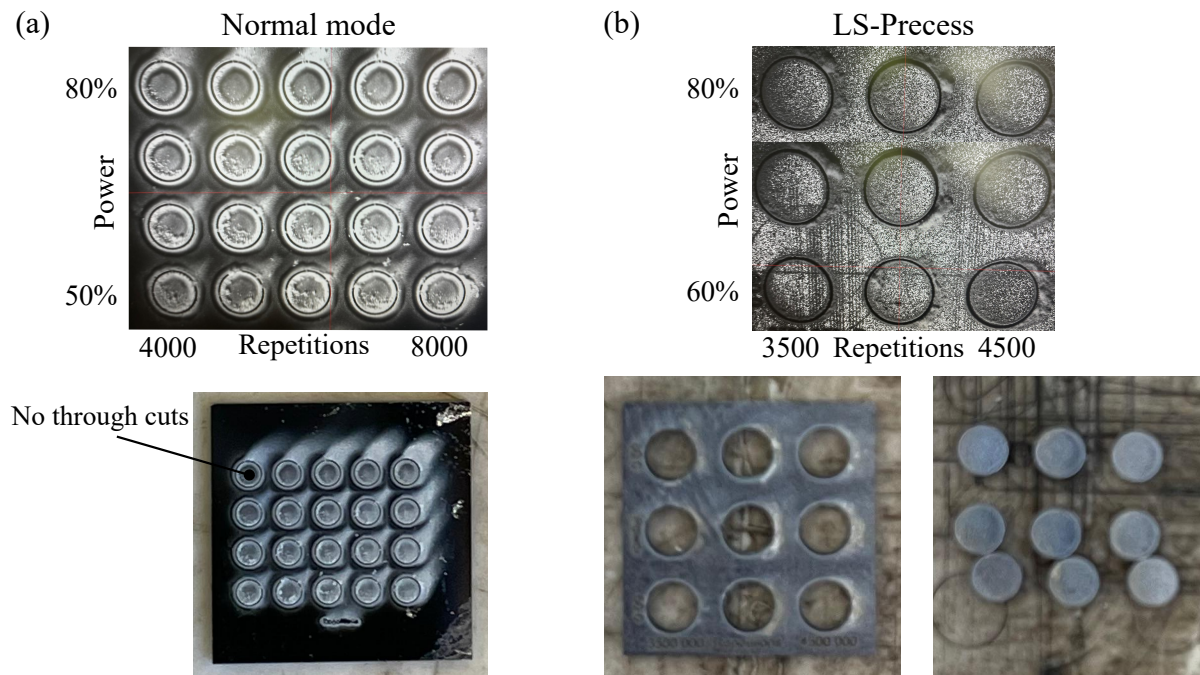


Figure 3.1: Comparison of femtosecond-laser through-cutting of 525 μm -thick silicon in normal and LS-Precess modes. (a) Normal mode: on-axis camera views of a power—repetition matrix (power 50 % to 80 %; repetitions 4000 to 8000). The photograph of the resulting die indicates no released features (no through-cuts) within this parameter window. (b) LS-Precess: matrix (power 60 % to 80 %; repetitions 3500 to 4500) with die and released-disk photographs confirming through-cuts at all tested points. Scan speed (200 mm/s), outlines (three), and outline pitch (0.01 mm) were constant.

3.3. PRODUCTION RECIPE, RELIABILITY–TIME TRADE-OFF, AND YIELD

After selecting LS-Precess as the enabling mode, the cutting parameters were refined to reduce machining time while preserving through-cut reliability and avoiding backside exit defects when removing the blades from the wafer. A key practical observation is that marginal or incomplete cuts can lead to macroscopic chipping at the kerf exit and local bulge formation where material remains partially attached before breaking away (Fig. 3.2(c)). Such defects are unacceptable for subsequent work because they introduce local geometric errors at the blade perimeter and can locally increase the realised radius.

Three process choices were made to balance reliability, throughput, and edge condition. First, the number of outlines was reduced after mode selection. While the mode-comparison matrices in Fig. 3.1 used three outlines for robustness, subsequent LS-Precess optimisation indicated that two outlines were sufficient to obtain reliable through-cuts, while reducing both kerf width and machining time. Second, the repetition count was set with a small safety margin rather than minimised aggressively, to mitigate occasional incomplete cutting. Third, the laser power was iterated within the policy-limited range and assessed qualitatively by edge appearance. Representative inspections suggested that higher power produced a straighter perimeter with fewer pronounced local notches in the near-edge region (Appendix C, Fig. C.1). On this basis, the final production recipe used 80 % power.

Using the resulting LS-Precess recipe (scan speed 200 mm s^{-1} , laser power 80 %, 1500 repetitions, two outlines with 0.01 mm outline pitch), cutting a single annular blade required approximately 42 min of machine time. With this recipe, four Si/a-SiC blades and four Si/SiO₂ blades were produced as usable annular disks without macroscopic backside chipping or bulge formation of the type shown in Fig. 3.2(c). These blanks formed the inputs for the subsequent underetching experiments (Challenge 3) and were later mounted for high-speed rotation experiments.

3.4. EDGE CONDITION AND IMPLICATION FOR MINIMUM UNDERETCH LENGTH

Although LS-Precess enables reliable through-cutting, the near-edge region of the laser-cut perimeter exhibits a laser-affected zone. Fig. 3.2 shows representative Keyence micrographs illustrating three relevant observations. First, when the functional film faces the laser ("film facing up"), a visibly affected near-edge




Figure 3.2: Representative edge-condition observations after femtosecond-laser cutting, illustrating near-edge damage and a critical failure mode. (a) Si/SiO₂ with the film facing the laser: a visibly affected near-edge region is present; annotated widths indicate representative local values (extent varies along the circumference). (b) Si/SiO₂ with the film facing away from the laser: the visibly affected region is reduced at the shown location, indicating a modest orientation benefit. (c) Example failure mode after an incomplete through-cut: backside chipping and a local bulge at the kerf exit, where partially attached material subsequently broke away. Scale bars as indicated.

region is present in which dimples, notches, and local film damage occur (Fig. 3.2(a)). The annotated widths indicate representative *local* values. Both the extent and morphology vary along the circumference. Second, when the film faces away from the laser ("film facing down"), the visibly affected region can be reduced at some locations (Fig. 3.2(b)). This orientation effect was recognised late in the project and is best interpreted as a modest improvement rather than a complete mitigation of near-edge damage. Third, incomplete cutting can produce severe backside exit defects (chipping and bulge formation), reinforcing that through-cut reliability is a geometric quality requirement rather than merely a convenience (Fig. 3.2(c)).

These observations motivate an explicit design rule for the subsequent underetching step. The laser-affected zone extends over length scales of order tens of micrometres from the nominal cut edge (representative local values are indicated in Fig. 3.2(a)). Forming a continuous ultra-thin cutting edge therefore benefits from an underetch length that *comfortably exceeds* this affected region beneath the future overhang. In this work, an underetch length of at least 60 μm was targeted to provide margin against local variability in near-edge damage and to increase the likelihood of achieving a continuous overhang along the full circumference.

3.5. BORE-FIT TUNING UNDER A HANDLING CONSTRAINT

In addition to perimeter integrity, the bore clearance should be minimised to reduce the potential for eccentric mounting, which would contribute directly to radial run-out during high-speed rotation (Challenge 1). In femtosecond-laser cutting, the realised cut edge is influenced by the finite kerf width, which depends on spot size and focus conditions, the selected cutting mode (normal versus LS-Precis), and process parameters, such that the final bore size cannot be inferred directly from the programmed radius alone. Fig. 3.3 provides




Figure 3.3: Supporting kerf-width and dimensional-bias characterisation for femtosecond-laser cutting in LS-Precess mode using simple test patterns. The finite kerf width causes the realised edges to deviate systematically from the programmed centreline, motivating empirical tuning of the programmed bore radius by stepwise fit tests on the shaft seat. Scale bars as indicated.

supporting evidence using a simple calibration pattern: two parallel lines with a nominal centre-to-centre spacing of 1 mm were cut in LS-Precess mode using two outlines (pitch 0.01 mm). Top-view Keyence measurements indicate a kerf width of $\approx 71 \mu\text{m}$ per line and a visibly affected zone of $\approx 25 \mu\text{m}$ adjacent to the kerf. The measured inner-edge spacing of $\approx 931 \mu\text{m}$ and outer-edge spacing of $\approx 1067 \mu\text{m}$ imply an effective edge offset of order half the kerf width (i.e. $\sim 35 \mu\text{m}$) relative to the programmed centreline. In this work, this bias was therefore treated empirically when tuning the programmed bore radius by stepwise fit tests on the shaft seat.

Accordingly, the bore was tuned by progressively reducing the programmed inner radius in small steps and evaluating the resulting fit on the shaft seat. The practical lower bound on clearance was determined not by a metrological target but by safe manual assembly. The brittle annular disks were mounted using a vacuum tweezer, which avoids contacting the outer edge yet limits the assembly force that can be applied. As the inner radius was reduced, the fit transitioned from visibly loose to tight but assemblable with the vacuum tweezer, and finally to "no fit" where further reduction would require excessive assembly force and risk chipping the inner edge. This stepwise approach yielded a repeatably mountable bore clearance while limiting perceptible lateral play. The practical implication of the single-point, compliant grip during seating is illustrated in Appendix C (Fig. C.2).

A clear direction for future work is to replace vacuum-tweezer mounting by a simple handling or pressing fixture that supports the disk and permits controlled axial assembly force without contacting the cutting edge. Such a fixture could allow smaller clearances to be explored safely, potentially reducing the scope for eccentric mounting and thereby further reducing radial run-out.

3.6. DISCUSSION AND IMPLICATIONS

Challenge 2 established a practical and repeatable route to fabricate annular blade blanks suitable for later experiments. LS-Precess was identified as the enabling intervention for through-cutting $525 \mu\text{m}$ silicon within practical machining times, whereas normal mode did not yield reliable through-cuts under the tested conditions. A production recipe was then selected that balances throughput (42 min per blade) against through-cut reliability, avoiding macroscopic backside exit defects that would compromise perimeter geometry. Keyence inspection further showed a laser-affected edge zone on length scales of order tens of micrometres, motivating a target underetch length of at least $60 \mu\text{m}$ to support continuous overhang formation in Challenge 3. Finally, bore-fit was tuned empirically by stepwise reduction of the programmed inner radius, with the practical limit set by safe manual assembly using a vacuum tweezer.

4

CHALLENGE 3: SELECTIVE UNDERETCH WITH HUB-BORE PROTECTION

The third challenge addressed in this work is to develop a fabrication strategy that releases the ultra-thin cutting edge while preserving the geometry of the hub bore (blade inner hole), in terms of both size and concentricity with the blade outer perimeter. As discussed in Chapter 2, the fit between the shaft blade seat and the blade hub bore influences the quasi-static radial run-out at the blade outer edge. Fig. 4.1(a) illustrates a diametral clearance, denoted here as s , between the shaft seat and the hub bore. This clearance permits an eccentric positioning of the blade relative to the shaft by up to $s/2$, thereby contributing up to s peak-to-peak quasi-static run-out at the blade outer edge. In addition, the concentricity between the hub bore and the blade outer perimeter directly affects run-out. If the centres of the inner and outer perimeters are offset by a distance d , Fig. 4.1(b), the resulting geometric run-out at the outer edge is $2d$ peak-to-peak. Consequently, the blade should be manufactured such that (i) the hub bore is minimised while still allowing manual assembly onto the shaft, avoiding damage to the brittle Si/thin-film structure (i.e., clearance fit), and (ii) eccentricity of the hub bore and outer perimeter is minimised.

The production of the final ultra-thin cutting blade comprises two main steps: (1) femtosecond-laser cutting of the blade outer perimeter and hub bore from the thin-film-coated wafer, and (2) selective plasma etching of the Si substrate at the blade outer perimeter to undercut the substrate and thereby release the thin-film overhang that acts as the cutting edge (SF_6 isotropic plasma etch; in the recipe used here, Si etches at approximately $12 \mu\text{m min}^{-1}$). This gives rise to two manufacturing routes.

Route 1: laser cut the outer perimeter first, perform the plasma etch to release the thin-film overhang, and then laser cut the hub bore. The principal advantage of this route is that the hub bore is created after the plasma step, so its Si sidewalls are not exposed to the etchant. This preserves the as-cut bore geometry

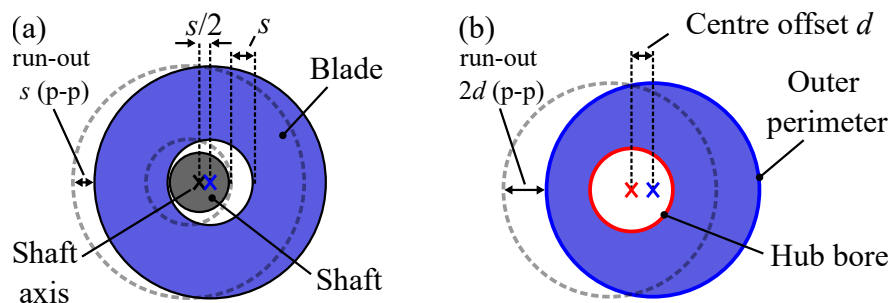


Figure 4.1: Schematic illustration of two geometric contributions to quasi-static radial run-out at the blade outer perimeter. Solid and dashed outlines indicate two opposite angular positions (180° apart). (a) A diametral clearance s between the shaft blade seat and the blade hub bore permits eccentric positioning of the blade relative to the shaft by up to $s/2$, which can contribute up to s peak-to-peak run-out at the outer perimeter. (b) Eccentricity between the centres of the hub bore and the outer perimeter by d introduces a geometric run-out of $2d$ peak-to-peak at the outer perimeter.

and therefore maintains the tightest achievable shaft–blade fit (minimum diametral clearance s), which minimises the associated quasi-static run-out contribution (Fig. 4.1(a)).

However, this route introduces two major drawbacks. First, the LS-Lab laser system requires manual selection of the pattern centre in the software using an on-axis camera view. Re-centring with respect to the already-cut outer perimeter is therefore expected to introduce additional and difficult-to-control eccentricity between the hub bore and outer perimeter. In principle, a centre marker could be produced during the first laser-cutting step, but in practice the on-axis view and user-interface do not provide sufficient means to reposition the centre repeatedly to that marker with the precision required to limit eccentricity to the tens-of-micrometres level.

Second, cutting the hub bore after the overhang has been released increases the risk of damage to the fragile cutting edge and complicates debris management during cutting. Additional handling steps increase the risk of accidental edge contact or dropping. Moreover, full-depth laser cutting generates debris that is normally mitigated using air assist and vacuum extraction while the sample is secured at its perimeter. After edge release, perimeter fixation becomes impossible due to the presence of the ultra-thin cutting edge. Reducing or disabling these measures, however, was observed to increase local debris accumulation along the cut path, which degrades cutting efficiency and dimensional accuracy.

Route 2: laser cut the complete annular geometry (outer perimeter and hub bore) in a single operation, followed by the plasma etch to release the cutting edge. This route has two key advantages. First, the inner and outer perimeters are generated from the same programmed geometry without any intermediate re-centring step. Hence, their mutual concentricity is primarily determined by the machine-controlled tool path definition and local process effects (e.g. kerf width, taper, and local process variations), rather than by a second manual alignment operation. Second, once the cutting edge is released, no subsequent fabrication steps are required, reducing handling and therefore lowering the probability of edge damage.

The disadvantage of Route 2 is that the hub bore already exists during the plasma etch. Without additional measures, the Si sidewalls of the hub bore would also be exposed to the etchant, increasing the bore diameter and thereby worsening the shaft–blade fit and the associated run-out. Although Route 2 therefore presents a clear risk to hub-bore integrity, it offers the greatest potential for low eccentricity between the hub bore and the outer perimeter, provided that a method is available to selectively protect the hub bore during the plasma step. This motivated the development of a dedicated masking and handling fixture, consisting of a *lid* and a *base*. The lid concept is introduced below. The base concept is introduced thereafter, before presenting the experimental evidence for their performance.

4.1. HUB-BORE PROTECTION DURING PLASMA ETCHING

Protecting the hub-bore sidewalls during the isotropic SF₆ plasma etch requires a solution that:

- selectively shields the hub bore while leaving the blade outer perimeter exposed for substrate under-cutting,
- is compatible with cleanroom processing and minimises contamination risk,
- can be applied and removed without contacting the blade outer edge, and
- remains in place during handling and rapid equipment stage motions.

These requirements were met by a reusable circular *lid* (Fig. 4.2). The lid is fabricated from a Si <100> wafer (nominal thickness 525 μm, p-type, single-side polished) and comprises a circular plate (outer radius 15 mm, thickness ≈ 345 μm) with a central cylindrical tower (radius 8.5 mm, height ≈ 180 μm), Fig. 4.2(a). During processing, the lid is placed on top of the annular blade with the tower facing downward and inserted into the hub bore, Fig. 4.2(b). This geometry provides two functions. First, it establishes a self-aligned sacrificial hard mask that physically blocks direct plasma exposure of the hub bore sidewalls, thereby suppressing etching of the bore during the SF₆ step. Second, the tower acts as a locating feature that constrains lateral motion of the lid relative to the blade, holding the lid in place by gravity alone (no surface contamination from adhesives or tapes).

Using Si for the lid ensures compatibility with standard cleanroom processing and avoids introducing metallic materials into the etch environment. Furthermore, the lid can be manipulated using a vacuum tweezer,




Figure 4.2: Lid concept for selective protection of the hub bore during isotropic SF_6 plasma etching. (a) Lid geometry: a stepped Si disk with outer radius 15 mm and a central locating tower (radius 8.5 mm) providing an approximately $180\ \mu\text{m}$ step height. (b) Cross-sectional schematic of the lid placed on the blade with the tower inserted into the hub bore, forming a self-aligned sacrificial hard mask that blocks plasma exposure of the hub bore sidewalls while leaving the outer perimeter exposed for substrate undercutting; the lid is held in place by gravity and can be handled using a vacuum tweezer. Arrows indicate incoming etch species and red shading indicates etch-exposed Si surfaces.

enabling placement and removal without contacting the blade outer edge (Fig. 4.2(b), Fig. 4.3). The lid is reusable for multiple etch cycles; in practice, reuse is limited by progressive consumption of the lid during repeated plasma exposure, which is addressed quantitatively in the results section.

4.1.1. LID FABRICATION

The lid is manufactured using the same LS-Lab femtosecond laser micromachining system used for blade cutting. In brief, the outer circular perimeter (radius 15 mm) is cut through the wafer thickness, and a circular pocket is subsequently removed over the annular area between 8.5 mm and 15 mm radius to form the raised central tower. Notably, this stepped geometry was realised using the laser system not only for through-cutting but also for depth-controlled pocket milling (hatch ablation), enabling the integrated locating feature without additional tooling or lithography.

An example of a fabricated lid is shown in Fig. 4.3. The resulting step height was verified using white-light interferometry, yielding a tower height of approximately $180\ \mu\text{m}$. The lid surface that contacts the blade is laser-machined (i.e. not polished), and therefore the lid-blade interface is not intended to be hermetic. Nevertheless, any potential leakage path from the exposed outer boundary to the shielded hub-bore sidewalls is long, and the effectiveness of the shielding is demonstrated experimentally in Section 4.3.

4.2. BLADE FIXTURING AND HANDLING FOR PLASMA ETCHING

The PlasmaPro 100 Estrelas system used for the SF_6 plasma etch in this work transfers wafers into the process chamber via a moving load lock. Consequently, small parts must be processed on a carrier wafer to ensure reliable transport and positioning. In standard plasma processing, a die is temporarily fixed to a carrier wafer using a thin film of vacuum oil, which provides capillary adhesion and improves thermal contact to the temperature-controlled carrier. After processing, the die can typically be removed by applying force to




Figure 4.3: As-fabricated Si lid produced by femtosecond-laser through-cutting and depth-controlled pocket milling. (a) Lid manipulated using a vacuum tweezer, enabling handling without contacting the outer edge. (b) Close-up view showing the raised central locating tower (radius 8.5 mm) and the laser-machined contact face (outer radius 15 mm) that interfaces with the blade during etching (see inset for orientation). The tower step height was verified by white-light interferometry ($\approx 180\ \mu\text{m}$).

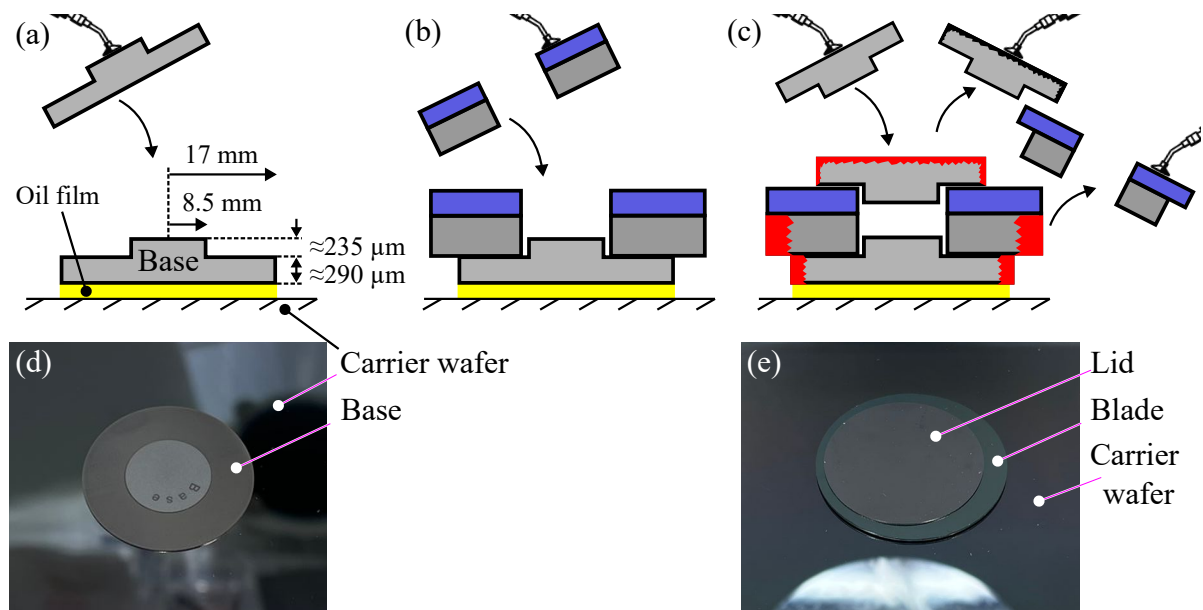


Figure 4.4: Base-assisted handling workflow for SF_6 plasma etching of the annular blade. (a) Base oil-bonded to carrier wafer (oil only at the base-carrier interface). (b) The blade is placed onto the base with the tower inserted into the hub bore, providing self-aligned lateral constraint while the blade is held by gravity. (c) Lid added to shield the hub-bore sidewalls while leaving the outer perimeter exposed (red: etch-exposed Si). (d) Photograph of the base positioned on the carrier wafer. (e) Photograph of the assembled stack (carrier wafer + base beneath, blade, and lid on top) prior to etching.

4

the its edge until it slides off the carrier wafer.

In the present work, however, this conventional approach becomes problematic once the ultra-thin cutting edge is released at the blade outer perimeter. Any direct contact with the outer perimeter (e.g. for perimeter fixation or manual removal) risks fracturing the fragile overhang. Removal by pushing on the hub-bore rim is also undesirable. During transfer from the carrier to a work surface, a small uncontrolled drop can nick the edge, and using tweezers near the hub bore introduces a risk of slipping and striking the cutting edge. These risks were considered unacceptable given the limited number of thin-film-coated wafers suitable for blade fabrication and the goal to minimise the probability of edge damage during processing and handling. Furthermore, directly oil-bonding the blade to the carrier wafer would necessitate solvent-based cleaning after processing, which was avoided to reduce handling and the risk of damage to the released edge.

These constraints motivated a dedicated support and handling solution with the following requirements:

- fix the blade during handling and rapid equipment stage motions without contacting the blade outer perimeter,
- allow placement and removal without contacting the blade outer perimeter, and
- be compatible with cleanroom processing while minimising contamination risk.

These requirements were met by a reusable circular *base* (Fig. 4.4), which is complementary to the *lid* introduced in Section 4.1. The base is fabricated from a Si $\langle 100 \rangle$ wafer (nominal thickness $525 \mu\text{m}$, p-type, single-side polished) and comprises a circular plate (outer radius 17 mm , thickness $\approx 290 \mu\text{m}$) with a central cylindrical tower (radius 8.5 mm , height $\approx 235 \mu\text{m}$), Fig. 4.4(a). The outer radius was intentionally chosen larger than that of the lid to (i) increase the support area beneath the blade away from the released edge region and (ii) maximise the conductive heat-transfer path from the blade through the base to the temperature-controlled carrier wafer. In addition, the base locally shields the blade underside in the supported region, thereby reducing unwanted backside etching and thickness loss in that area.

During processing, the base is fixed to the carrier wafer using a controlled thin film of vacuum oil at the *base-carrier* interface only (i.e. no adhesives, tapes, or oil are applied to the blade itself). The blade is then placed on the base with the tower facing upward and inserted into the hub bore, Fig. 4.4(b). The tower acts as a self-aligning locating feature that constrains lateral motion of the blade relative to the base, while the blade is held

against the base by gravity alone. The lid is subsequently placed on top of the blade as described previously (Fig. 4.4(c)), so that the hub-bore sidewalls are shielded from above while the blade outer perimeter remains exposed for substrate undercutting.

After etching, the lid is first removed using a vacuum tweezer. The blade can then be lifted off the base without contacting the outer perimeter because it is not bonded to the base. Finally, the base can be removed from the carrier wafer using standard tweezers, as it is sacrificial and oil-bonded only at the base-carrier interface.

4.2.1. BASE FABRICATION AND CONTACT-SURFACE TOPOGRAPHY

The base is fabricated from the same Si wafer material and using the same femtosecond-laser process settings as the lid, but used in the inverted orientation. Consequently, the base contact surface is representative of the lid-blade contact surface as well. White-light interferometry over an area of approximately $0.31 \text{ mm} \times 0.23 \text{ mm}$ (after tilt removal) shows micron-scale surface topography on the laser-machined contact surface. Representative line profiles indicate a peak-to-valley height variation of approximately $14 \mu\text{m}$ to $17 \mu\text{m}$, excluding isolated invalid points associated with local measurement dropouts. This confirms that the interface is non-hermetic and may contain micro-channels, which is considered in the interpretation of the masking performance in Section 4.3.

4.2.2. SELECTION OF TARGET UNDERCUT (OVERHANG) LENGTH

The isotropic SF_6 plasma step defines the radial undercut length of the Si substrate and therefore the released thin-film overhang length L that forms the cutting edge. The target range for L was selected based on two practical constraints.

First, femtosecond-laser cutting introduces a near-edge region in the thin film, characterised by local defects (e.g. dimples, ridges, and occasional chipping) extending laterally from the nominal kerf in the order of tens of micrometres (Section 3.4). To ensure that the released overhang forms a continuous edge around the full circumference, the undercut should exceed this affected zone. Based on laser-cut inspections, an undercut of at least $50 \mu\text{m}$ to $60 \mu\text{m}$ is therefore required. In practice, to provide margin for process variability while still aiming for a continuous edge, the lowest targeted overhang length was set to $L = 60 \mu\text{m}$.

Second, the released thin-film edge may exhibit out-of-plane curvature after release due to residual stress gradients in the film [44], [50], [51]. To obtain an order-of-magnitude estimate, the edge was modelled locally as a prismatic Euler-Bernoulli cantilever beam of length L and film thickness d , for which the initial tip deflection caused by a through-thickness linear stress gradient $\Delta\sigma/d$ can be approximated as [44], [52]

$$h_0 \approx \frac{\Delta\sigma}{d} \frac{L^2}{2E}, \quad (4.1)$$

highlighting the strong length dependence. In this work, the stress gradient is not known in advance (conventionally it is determined experimentally from the curvature of released test cantilevers) and is highly sensitive to deposition conditions [44]. Therefore, a parametric sweep was performed over a literature-informed range




Figure 4.5: Centrifugally stiffened cantilever-beam model of the ultra-thin cutting edge. (a) Schematic of the rotating dicing blade, indicating the inner and outer blade radii r_i and r_o , the clamping radius R_E of the overhanging cutting edge, and the direction of rotation Ω . (b) The cutting edge is idealised as a slender cantilever of length L , in-plane width w , and thickness d , attached to the blade body at radius R_E . Local beam coordinate system with axial coordinate x measured from the clamped end ($x = 0$) to the free tip ($x = L$) and transverse deflection $w(x)$ in the z -direction, as used in the Rayleigh-Ritz analysis of centrifugal stiffening.

of $\Delta\sigma/d = 50\text{MPa}\mu\text{m}^{-1}$ to $300\text{MPa}\mu\text{m}^{-1}$ [44], [51], [53], [54]. For a 96 nm a-SiC film, this implies that beams with $L \lesssim 40\mu\text{m}$ remain nearly straight ($h_0/L \lesssim 3\%$), whereas the most extreme combinations (e.g. $L = 150\mu\text{m}$ at the highest stress gradients) can reach curvature on the order of $h_0/L \approx 10\%$ (Appendix D).

During operation, rotation introduces centrifugal axial tension in the overhang, which adds geometric stiffness and reduces the initial deflection. This speed-induced stiffening is exploited in ultra-thin dicing blades, where increasing rotational speed helps suppress blade bending [28], [29]. To estimate this effect, the cutting edge is again idealised as a centrifugally stiffened cantilever (Fig. 4.5). Using a simple one-mode Rayleigh–Ritz (energy) formulation [52], the reduced tip deflection is

$$\frac{h(\Omega)}{h_0} = \frac{1}{1 + \alpha}, \quad \alpha = \frac{\rho A \Omega^2 L^3 (27L + 35R_E)}{280EI}, \quad (4.2)$$

with $R_E = r_o - L$, $A = wd$, and $I = wd^3/12$. Equivalently, the stiffening factor follows as $k_{\text{eff}}/k_0 = 1 + \alpha$ (Appendix E).

Evaluated at $N = 20000$ RPM as a design-study benchmark, this model predicts negligible stiffening for $L \lesssim 40\mu\text{m}$ ($k_{\text{eff}}/k_0 \approx 1.00\text{--}1.01$, $<2\%$ deflection reduction), modest stiffening for $L = 60\mu\text{m}$ to $80\mu\text{m}$ (roughly 4% to 9% stiffness increase), and substantial stiffening for $L = 100\mu\text{m}$ to $150\mu\text{m}$ ($k_{\text{eff}}/k_0 \approx 1.2\text{--}1.6$, corresponding to 15% to 40% deflection reduction). These trends imply a trade-off: short overhangs are intrinsically stiff and nearly flat but gain little additional stiffness from rotation, whereas longer overhangs may exhibit larger initial curvature (depending on the uncertain stress gradient) yet stiffen substantially during operation. Considering that spindle vibrations are on the order of several tens of micrometres in the present setup, a moderate level of compliance was considered desirable to reduce the likelihood of brittle fracture under unavoidable vibration induced bending. Consequently, the fabrication focused on nominal overhang lengths in the range $60\mu\text{m}$ to $150\mu\text{m}$.

Finally, although the original intent was to etch Si/a-SiC annular blades, their availability was limited (only one Si(100) wafer with a 96 nm a-SiC film was available) and initial trials revealed unexpected behaviour of the a-SiC layer during the SF_6 process (Section 4.3.2). Subsequent etching experiments therefore used a readily available Si/SiO₂ wafer from the university inventory (approximately $4\mu\text{m}$ SiO₂ on the polished side and $10\mu\text{m}$ on the unpolished side). Repeating the same stiffening estimate for this thicker $4\mu\text{m}$ SiO₂ edge clarifies how strongly the conclusions depend on film thickness. Since $\alpha \propto A/I \propto 1/d^2$, increasing thickness from 96 nm to $4\mu\text{m}$ suppresses centrifugal stiffening by $\sim (4/0.096)^2$, so that even at very high speeds (e.g. 60000 RPM) the predicted deflection reduction remains $\lesssim 1\%$ for $L = 150\mu\text{m}$. In other words, rotational stiffening is an important mechanism for ultra-thin film edges, but becomes negligible for the much thicker SiO₂ case, where the edge is intrinsically stiff. In principle, α could be increased by enlarging the undercut length L (since α grows strongly with L), but compensating the $\sim (4/0.096)^2$ thickness penalty would require L to increase by about $\sqrt[3]{(4/0.096)^2} \approx 12$, i.e. from $150\mu\text{m}$ to the millimetre scale. Such large undercuts were not pursued to limit risk and because the study aimed to characterise the practically relevant range of undercuts with minimal sample consumption.

4.3. ETCHING RESULTS AND FIXTURE PERFORMANCE

The lid–base fixture was evaluated experimentally by etching annular blades in an isotropic SF_6 plasma process. The process calibration used in this work corresponds to a nominal Si removal of approximately $12\mu\text{m}$ per 60 s. The targeted outcome was the controlled release of thin-film overhang lengths in the range $L = 60\mu\text{m}$ to $150\mu\text{m}$ while maintaining hub-bore integrity.

4.3.1. SUMMARY OF ETCH TRIALS

Table 4.1 summarises the etch attempts and the primary outcomes. The initial trials were performed on Si/a-SiC blades (nominal a-SiC thickness 96 nm), motivated by the expected high selectivity of a-SiC in SF_6 . Due to unexpected degradation of the a-SiC layer during these trials (Section 4.3.2), subsequent process-development experiments were performed on an inventory Si/SiO₂ wafer (approximately $4\mu\text{m}$ oxide on the polished side and $10\mu\text{m}$ on the unpolished side), which enabled systematic evaluation of undercut lengths and fixture sealing behaviour (Section 4.3.3).

Table 4.1: Summary of isotropic SF₆ plasma etch trials performed with the lid–base fixture. Undercut lengths are nominal values based on the process calibration (approximately 12 μm per 60 s) unless stated otherwise.

| Blade system | Attempt | Time | Target L | Observed outcome (high level) |
|---------------------|---------|---------|------------|--|
| Si/a-SiC | Etch 1 | 5 min | 60 μm | a-SiC removed outside lid; no overhang observed |
| Si/a-SiC | Etch 2 | 1 min | 12 μm | partial survival; overhang of 11 μm to 12 μm observed |
| Si/a-SiC | Etch 3 | 2 min | 24 μm | a-SiC removed; no overhang observed |
| Si/a-SiC | Etch 4 | — | — | alternative recipe; a-SiC removed |
| Si/SiO ₂ | Etch 1 | 5 min | ≈ 60 μm | released SiO ₂ overhang; lid intact |
| Si/SiO ₂ | Etch 2 | 7.5 min | ≈ 100 μm | released SiO ₂ overhang; lid intact |
| Si/SiO ₂ | Etch 3 | 10 min | ≈ 120 μm | released SiO ₂ overhang; lid consumed late in cycle |

4.3.2. ETCHING RESULTS FOR Si/A-SiC BLADES

The Si/a-SiC etch trials are reported first because they represent the intended blade concept (ultra-thin a-SiC edge). The results are presented chronologically to separate evidence for lid-defined selectivity on the top surface from the unexpectedly limited survival of the a-SiC film in the exposed region.

ETCH 1: 5 min EXPOSURE (TARGET $L \approx 60 \mu\text{m}$)

A first Si/a-SiC blade was processed for 5 min, corresponding to a nominal undercut of approximately 60 μm. Prior to etching, the a-SiC-coated surface exhibited a clear glossy white (green when observed by eye) appearance under the microscope (Fig. 4.6(a)). After the etch, the exposed annular region outside the lid footprint changed to a matte grey appearance consistent with the underlying Si substrate (Fig. 4.6(b)). The transition area (Fig. 4.6(c)) indicates that the lid effectively suppressed plasma attack on the top surface in the covered region.

Microscopy of the exposed region showed surface morphology consistent with loss of the a-SiC layer (rough Si-like texture in Fig. 4.6(d)), and no released a-SiC overhang was observed at the outer perimeter. Together, these observations indicate that the a-SiC layer was removed (or thinned below functional continuity) under these conditions, despite the expected selectivity of a-SiC in SF₆.

ETCH 2: 1 min EXPOSURE (TARGET $L \approx 12 \mu\text{m}$)

To reduce the plasma exposure, a second Si/a-SiC blade was etched for 1 min (nominal undercut 12 μm). In comparison to Etch 1, the post-etch surface retained visibly more of the original a-SiC appearance outside the protected area of the lid (Fig. 4.7(a)), suggesting partial survival of the film. However, microscopy still showed clear evidence of film degradation outside the lid footprint (colour change of the a-SiC film, Fig. 4.7(b)), implying that the a-SiC thickness was reduced during the etch.

At the outer perimeter, a released overhang length of 11 μm to 12 μm was observed in microscopy (Fig. 4.7(c)). The edge was not continuous around the full circumference with local regions showing interruptions due to edge damage originating from the femtosecond-laser cutting. Despite these non-uniformities, the presence of a released, sub-100 nm-class film overhang demonstrates that the ultra-thin edge can remain mechanically self-supporting at short released lengths and does not collapse under its own weight.




Figure 4.6: Optical microscopy of a Si/a-SiC blade before and after Etch 1 (SF₆, 5 min; target undercut $L \approx 60 \mu\text{m}$). (a) Pre-etch a-SiC surface. (b) Post-etch blade showing a matte-grey exposed annulus outside the lid footprint. (c) Close-up of the protected–exposed boundary demonstrating a sharp lid-defined transition. (d) Rough Si-like morphology in the exposed region consistent with a-SiC film loss. Red markers indicate magnified locations. Scale bars: 1 mm (a–b), 100 μm (c–d).




Figure 4.7: Optical microscopy of the second Si/a-SiC blade after Etch 2 (SF_6 , 1 min; nominal undercut target $L \approx 12 \mu\text{m}$). (a) Post-etch overview showing the lid-defined protected footprint and the exposed annulus outside the lid. The red box indicates the region magnified in (b). (b) Close-up of the protected-exposed transition, where a distinct contrast change indicates partial degradation of the a-SiC film in the exposed region while the protected area retains its original appearance. (c) Outer perimeter microscopy showing a released a-SiC cutting-edge overhang with length $L \approx 11 \mu\text{m}$. Scale bars: 1 mm (a), 100 μm (b), and 50 μm (c).

ADDITIONAL Si/a-SiC ATTEMPTS AND INTERPRETATION

A third Si/a-SiC blade etched for 2 min again resulted in removal of the a-SiC layer, and a fourth trial using an alternative recipe also resulted in full loss of the film. No direct wafer-temperature data were obtained during these experiments. However, a plausible explanation for the unexpected degradation of the a-SiC films is insufficient heat extraction from the blade through the base to the temperature-controlled carrier wafer, leading to elevated local temperature and degraded selectivity for a-SiC.

Importantly, across the Si/a-SiC trials the lid consistently produced a sharply delineated protected footprint on the top surface (Fig. 4.6(b), Fig. 4.7(a)), demonstrating that the lid-base concept can achieve strong spatial selectivity even when the lid-blade interface is non-hermetic (Section 4.2.1).

4.3.3. ETCHING RESULTS FOR Si/SiO₂ BLADES AND HUB-BORE PROTECTION

The Si/SiO₂ experiments are used here as process-development trials to (i) evaluate achievable undercut lengths with minimal sample consumption and (ii) quantify the hub-bore shielding performance of the lid-base fixture. The thicker oxide films are not treated as performance-representative substitutes for the a-SiC edge. Rather, they enable robust inspection of fixture behaviour and released-edge geometry.

RELEASED OVERHANGS AND LID ENDURANCE

Three Si/SiO₂ blades were etched with nominal targets of $L \approx 60 \mu\text{m}$, $100 \mu\text{m}$ and $120 \mu\text{m}$ (Table 4.1). In all three cases, microscopy confirmed a released SiO₂ rim at the outer perimeter (Fig. 4.8). On the polished side ($t \approx 4 \mu\text{m}$), the released overhang length measured along the perimeter was $L = 67.6 \pm 5.4 \mu\text{m}$, $97.4 \pm 4.3 \mu\text{m}$, and $121.4 \pm 3.8 \mu\text{m}$ for the nominal $60 \mu\text{m}$, $100 \mu\text{m}$ and $120 \mu\text{m}$ blades, respectively (each based on $n = 52$ local measurements distributed around the perimeter; reported as mean \pm standard deviation). Compared with the Si/a-SiC trials, these experiments provided stable process-development data because the oxide film is substantially thicker and remained present throughout the etches required to reach the targeted undercuts.

In addition to the front-side rim, the thicker rear-side SiO₂ film ($t \approx 10 \mu\text{m}$) formed a longer released "back-side" rim. The corresponding measured rear-side lengths were $L = 72.6 \pm 2.4 \mu\text{m}$, $105.3 \pm 3.1 \mu\text{m}$, and $137.9 \pm 2.8 \mu\text{m}$ for the nominal $60 \mu\text{m}$, $100 \mu\text{m}$ and $120 \mu\text{m}$ blades, respectively (again $n = 52$). The origin of this length asymmetry is not fully resolved, but three plausible (and not mutually exclusive) contributors are: (i) a pre-existing geometric difference between the front and rear edges introduced during femtosecond laser cutting (e.g. taper or different kerf geometry on the two surfaces), such that the apparent front- and rear-side protrusions differ already before etching; (ii) side-dependent effective plasma exposure during processing, which may alter the degree of erosion at the free oxide rim; and (iii) some recession of the free oxide rim during SF_6 exposure (i.e. non-zero oxide erosion), for which the thicker rear-side oxide would remain intact for longer than the thinner ($\sim 4 \mu\text{m}$) front-side film, leading to a slightly longer remaining backside rim.

The lid is a consumable mask, as cumulative plasma exposure progressively etches the lid surfaces. In the SiO₂ series, the annular shielding portion of the lid (contact-face region) was reported by the operator to be absent upon removal after the 10 min etch of the $120 \mu\text{m}$ edge (tower still present), implying that the hub bore was exposed during the final part of that cycle. This event provides a useful failure-case datapoint for interpreting hub-bore protection performance.




Figure 4.8: Optical microscopy of released SiO₂ cutting-edge rims at the outer perimeter of three Si/SiO₂ blades etched with nominal undercut targets of $L \approx 60\mu\text{m}$, $100\mu\text{m}$ and $120\mu\text{m}$. (a–c) Representative edge micrographs showing a continuous released oxide overhang adjacent to the Si blade body. The measured released lengths are $L = 65\mu\text{m}$, $99\mu\text{m}$, and $120\mu\text{m}$, respectively (red dimension arrows). Scale bars: $50\mu\text{m}$.

QUANTITATIVE BORE-RADIUS EVIDENCE FROM KEYENCE MEASUREMENTS

Hub-bore preservation was assessed using Keyence VHX-6000 stitched imaging at X200 magnification over an area of approximately $21\text{ mm} \times 21\text{ mm}$. For each blade, the inner perimeter was inspected in the stitched view and the radius r_i was estimated by selecting three perimeter points along the perimeter, from which the software returned a circle-fit radius. This procedure was repeated $n = 10$ times with selections distributed around the perimeter (i.e. different angular locations), providing an empirical estimate of measurement scatter.

For the $L \approx 60\mu\text{m}$ SiO₂ blade, the post-etch bore radius was $r_i = (10010.3 \pm 6.9)\mu\text{m}$ (mean \pm standard deviation, $n = 10$) compared to a pre-etch value of approximately $r_i \approx 10008\mu\text{m}$. For the $L \approx 100\mu\text{m}$ blade, the post-etch radius was $r_i = (10008.6 \pm 9.8)\mu\text{m}$ ($n = 10$) compared to a pre-etch value of approximately $r_i \approx 10010\mu\text{m}$. In both cases, the observed difference between pre- and post-etch radii is within the measurement scatter, indicating no measurable bore enlargement. Consistent with this, optical inspection of the top-surface rim near the bore edge of the post-etch intact-lid case (Fig. 4.9(b)) showed no obvious etch attack when compared to a pre-etch SiO₂ image (Fig. 4.9(a)).

For the $L \approx 120\mu\text{m}$ blade, where lid loss was reported late in the etch cycle, the post-etch bore radius increased to $r_i = (10015.5 \pm 13.5)\mu\text{m}$ ($n = 10$) relative to a pre-etch value of approximately $r_i \approx 10010\mu\text{m}$. The mean shift and increased scatter is consistent with a reduction in circularity and/or localised perimeter changes. Correspondingly, microscopy of the top-surface rim near the bore edge showed clear morphological changes consistent with etch exposure (Fig. 4.9(c)). Together, these observations support the interpretation that the lid provided effective hub-bore shielding during the earlier etches SiO₂ etches (and the Si/a-SiC etches), while late-cycle lid consumption in the longest etch resulted in measurable bore enlargement.

Only the top surface and the top-surface rim at the bore edge were imaged. No direct sidewall imaging of the hub bore was obtained. Therefore, the demonstrated protection pertains to preservation of the fitted bore radius (within measurement scatter) and the absence of visible etch attack on the accessible top-surface features.

4.3.4. RELEASED-EDGE CURVATURE AND STRESS-GRADIENT CONSISTENCY CHECK (SiO₂)

The released SiO₂ edges exhibited an upward curvature after etching. To provide a compact link back to the stress-gradient discussion (Section 4.2.2), the initial out-of-plane deflection was estimated by focusing the Keyence microscope on (i) the base surface and (ii) the tip of the released edge at X2000 magnification. The difference in focus height was used as an estimate of the tip deflection h_0 . Treating the released rim locally as a cantilever beam (as in (4.1)) yields the corresponding order-of-magnitude stress-gradient estimate

$$\frac{\Delta\sigma}{d} \approx \frac{2Eh_0}{L^2}, \quad (4.3)$$

where $E \approx 70\text{ GPa}$ was used for SiO₂.

Using $L \approx 67.6\mu\text{m}$ and $h_0 \approx 1.13\mu\text{m}$ gives $\Delta\sigma/d \approx 35\text{ MPa}/\mu\text{m}$; for $L \approx 97.4\mu\text{m}$ and $h_0 \approx 4.14\mu\text{m}$ gives $\Delta\sigma/d \approx 61\text{ MPa}/\mu\text{m}$; and for $L \approx 121.4\mu\text{m}$ and $h_0 \approx 9.43\mu\text{m}$ gives $\Delta\sigma/d \approx 90\text{ MPa}/\mu\text{m}$. These values are intended as a consistency check rather than a precise material characterisation: (i) the rim geometry differs from an

ideal cantilever, and (ii) both L and h_0 are subject to measurement uncertainty. Nevertheless, the estimates indicate stress-gradient magnitudes of the same order for the shorter edges, while the longest edge shows a higher inferred gradient, which may reflect geometric/model-form differences or local process effects.

4.4. DISCUSSION AND IMPLICATIONS

The results demonstrate that the lid–base fixture provides strong selectivity during isotropic SF_6 etching. In both Si/a-SiC and Si/SiO₂ trials, post-etch inspection consistently revealed a sharply delineated boundary between the lid-covered region and the exposed region, indicating that the lid suppresses plasma attack on features within its footprint. For Si/SiO₂ blades etched with the lid intact (nominal $L \approx 60 \mu\text{m}$ and $100 \mu\text{m}$), Keyence-derived circle-fit estimates of the hub-bore radius remained unchanged within measurement scatter, and bore-edge microscopy showed no obvious rim degradation when compared to the pre-etch baseline. Together, these observations support the conclusion that the fixture preserves hub-bore geometry to within the resolution of the present metrology, while acknowledging that no bore-sidewall imaging was obtained.

A further strength of the approach is that lid consumption provides a practical failure indicator. In the longest Si/SiO₂ etch ($L \approx 120 \mu\text{m}$), the annular shielding portion of the lid was reported by the operator to be absent upon removal, implying late-cycle exposure of the hub region. In this case, the circle-fit radius estimates exhibited increased scatter and the bore-edge rim showed clear morphological changes consistent with etch attack. This failure case links loss of masking to observable hub-bore degradation, thereby strengthening the interpretation of the intact-lid results.

In contrast, reliable formation of long ultra-thin a-SiC edges was not achieved with the present process configuration. Etch 1 (5 min) and Etch 3 (2 min) resulted in removal of the a-SiC film in the exposed region, whereas Etch 2 (1 min) produced a short released overhang ($L \approx 11 \mu\text{m}$ to $12 \mu\text{m}$) that remained mechanically self-supporting but was not circumferentially continuous. Since no wafer-temperature data were available, the proposed explanation is hypothesis-driven: insufficient heat extraction through the base to the temperature-controlled carrier may elevate local temperature and reduce the effective selectivity for a-SiC. This suggests that future iterations should prioritise thermal management (e.g. improved conductive coupling and/or reduced thermal resistance in the stack) and/or revised etch conditions to widen the selective process window for ultra-thin a-SiC.

Overall, Challenge 3 establishes the lid–base fixture as a cleanroom-compatible, adhesive-free method for selective hub protection and safe handling of fragile released edges. The principal limitation observed in this work is not the masking concept itself, but the robustness of selective undercutting beneath ultra-thin a-SiC in the current thermal/process configuration.




Figure 4.9: Optical microscopy of the top-surface rim near the hub bore edge for Si/SiO₂ blades, illustrating hub-bore protection performance of the lid. (a) Representative pre-etch bore-edge morphology (baseline, prior to all SiO₂ etches). (b) Post-etch bore edge for a blade etched to $L \approx 60 \mu\text{m}$ (representative also for $L \approx 100 \mu\text{m}$), showing no obvious etch attack at the accessible top-surface rim. (c) Post-etch bore edge for the $L \approx 120 \mu\text{m}$ blade, for which lid loss was reported in the etch cycle, showing clear morphological change consistent with etch exposure at the rim. Insets provide magnified views of the indicated regions. Scale bars: 25 μm .

5

CHALLENGE 4: INTEGRATED PROOF-OF-CONCEPT CUT DEMONSTRATION

The final challenge is an integrated proof-of-concept demonstration of the released blade under realistic operating conditions: rotation to the minimum cutting-edge speed, survival of the released rims at speed, and controlled engagement with a sample. Following initial tests on a compliant sample (thread), the chapter centres on a controlled rigid-sample trial that preserves a measurable kerf signature and enables quantitative assessment of kerf geometry.

5.1. EXPERIMENTAL CONSTRAINTS AND APPROACH

Executing a first rigid cut with an ultra-thin released rim presents two primary risks: (i) over-travel leading to catastrophic fracture or substrate sidewall contact, and (ii) insufficient observability of first engagement. A stepwise approach protocol was therefore adopted to mitigate loading while retaining optical access.

A camera (Sony A7III with Laowa 24 mm probe lens) was aligned approximately perpendicular to the blade surface (axial viewing direction, Fig. 5.1(a)) to provide qualitative clearance indication. At rim speeds of 30 m/s ($r = 18$ mm), the edge cannot be frozen under continuous illumination at practical exposure settings and therefore appears motion blurred. Imaging therefore provides only approximate clearance information and cannot reliably localise first contact. Strobed illumination was evaluated, but rolling-shutter artefacts and lack of phase-lock limited its usefulness. Representative imaging trials (stationary, rotating, fast shutter, and strobed) are summarised in Appendix G.

The sample was mounted on a motorised linear stage (Thorlabs MTS25(/M)-Z8) providing a minimum repeatable increment of 0.8 μ m. Engagement was performed incrementally with frequent retraction and inspection to minimise the risk of uncontrolled penetration. The spindle-stage assembly is shown in Fig. 5.1. For clarity, the configuration shown uses the thread sample holder (Tests 2–4). An analogous holder was used for the rigid LEGO trials, with the sample fixed by tape.

5.2. SURVIVAL AND COMPLIANT ENGAGEMENT TESTS

The following tests were performed using a released Si/SiO₂ blade with nominal protrusion length $L \approx 120$ μ m. To preserve the axial camera view during approach, the thread sample must be held without obscuring the line-of-sight to the cutting edge. A 3D-printed sample holder was therefore designed to clamp the sample in a diagonal orientation between two arms (see Appendix G, Fig. G.2). This geometry keeps the region of interest visible when the sample is brought close to the blade (Fig. 5.1(b)).

5.2.1. TEST 1: SURVIVAL AT SPEED

The released SiO₂ blade was accelerated to approximately 17.6 kRPM and held for a short duration. No macroscopic breakaway of either the front- or rear-side edge was observed, confirming mechanical survival of the released edges under centrifugal loading and operational displacement levels established in Challenge 1.




Figure 5.1: Overview of the approach configuration used for the cut demonstrations, showing how optical access and controlled sample translation were achieved. (a) Top view indicating the axial camera placement relative to the blade rim, enabling visual clearance indication during the stepwise approach. (b) Front view of the blade–sample interface, illustrating the unobstructed line-of-sight to the primary contact zone at the cutting edge. (c) Isometric view of the integrated assembly (spindle and clamped blade, diagonal sample holder, and motorised linear stage), with the sample position indicated in green.

5.2.2. TESTS 2–4: CUTTING OF A TENSIONED THREAD

To establish whether the released edges exhibit cutting action while reducing the risk of immediate catastrophic failure, a compliant sample was selected: polyester sewing thread of diameter ≈ 0.1 mm. The thread also served as a practical contact indicator: under tension, fibre separation upon severing provides a clear qualitative signature of engagement without requiring a sharply resolved view of the edge. Partial (Test 2 and Test 3) and full severing (Test 4) were achieved at ≈ 16.2 kRPM without catastrophic edge failure. These tests confirm that at least one protruding edge can perform functional engagement under compliant conditions. However, the thread does not preserve a diagnostic surface track and does not allow attribution to either the ~ 4 μ m front-side rim or the ~ 10 μ m rear-side edge.

Detailed post-test microscopy of the edges is provided in Appendix G, Section G.4. In particular, the rear-side edge exhibited spatially localised slight wear in some regions while other inspected regions appeared unchanged, consistent with intermittent contact occurring over a limited circumferential sector. Example thread conditions before and after partial cutting are shown in Appendix G, Fig. G.3.

5.3. RIGID-SAMPLE TESTS

Rigid-sample testing was used to preserve a surface track for post-test microscopy and to probe which edge engages. Two LEGO (ABS) trials were performed: an initial over-travel case (Test 5) that exposes failure modes of the current observability strategy, followed by a controlled incremental approach (Test 6) that yields a continuous kerf and enables quantitative depth assessment.

5.3.1. TEST 5: OVER-TRAVEL AND SUBSTRATE SIDEWALL CONTACT

An initial rigid-sample trial was performed on a white LEGO (ABS) brick at $N \approx 16.9$ krPM. The sample mounting compromised the axial camera view and prevented a clear view of the surface immediately prior to contact. As a consequence, detection of first contact during rotation remained unreliable.

The sample was advanced in minimum-repeatable steps of $0.8\mu\text{m}$ with pauses between steps. Three fine-approach sequences of approximately $100\mu\text{m}$ each were executed, with retraction and visual inspection between sequences. During the final sequence, a subtle change in the apparent edge region was noted in the camera view (retrospectively consistent with the onset of contact), but the approach was continued.

Post-test inspection revealed a wide damage band of $w \approx 547\mu\text{m}$, comparable to the total blade stack thickness ($\sim 539\mu\text{m}$), consistent with over-travel leading to substrate sidewall involvement rather than cutting dominated by the released rim(s) (Fig. 5.2). The front-side rim was lost circumferentially, with only small residual chips remaining. The rear-side rim was removed over a large sector (on the order of more than half the circumference), again leaving only small residual chips in the affected regions. In the remaining sector(s), the rear-side rim showed no visible degradation and its protrusion length was indistinguishable from the pre-test condition within the inspection resolution. The full micrographs are provided in Appendix G (Fig. G.5). Overall, Test 5 demonstrates that, without a reliable first-contact indicator and with limited optical access, a small additional advance can transition from initial engagement to substantial over-travel and severe rim loss.

5.3.2. TEST 6: CONTROLLED INCREMENTAL RIGID CUT

A second rigid-sample test was performed on the same LEGO (ABS) sample using the blade with nominal protrusion length $L \approx 100\mu\text{m}$ (front-side thickness $\sim 4\mu\text{m}$, rear-side thickness $\sim 10\mu\text{m}$). The sample was advanced in increments of approximately $15\mu\text{m}$ (steps size of $1\mu\text{m}$), with full retraction and visual inspection of the surface and a camera check of the edge region after each increment. While demounting the sample for microscopy after each increment would provide higher sensitivity to the onset of a minimal kerf, remounting would introduce re-positioning errors which would negate the stage-cutting edge reference positions, making between-step microscopy impractical for this test. First visible engagement occurred after 13 increments at an operating speed of approximately 17.7krPM. No additional advance beyond initial engagement was performed.

Kerf morphology. Optical microscopy of the sample surface reveals a single continuous kerf of length ≈ 1.5 mm (Fig. 5.3(a)). The kerf width measured near the centre was approximately $46\mu\text{m}$, while both far ends of the kerf start with a narrow track of order $5\mu\text{m}$ and widen towards the middle (Fig. 5.3(b,c)). One plausible explanation for this systematic widening is a longer effective contact duration near the mid-length region, combined with operational displacement that causes the engaged rim to sweep laterally across the surface during contact. No wide damage band indicative of substrate sidewall contact was observed. The morphology is therefore consistent with rim-dominated kerf formation.

Depth measurement and aspect ratios. White-light interferometry was used to determine the maximum trench depth (peak-to-valley) relative to the local mean surface plane outside the kerf (representative map and line profile in Appendix G, Fig. G.6). Across five measurements near the kerf centre, the maximum trench depth ranged from $11.72\mu\text{m}$ to $14.71\mu\text{m}$ (mean $\approx 13.89\mu\text{m}$), while the corresponding kerf width ranged from $43.41\mu\text{m}$ to $47.02\mu\text{m}$ (mean $\approx 44.59\mu\text{m}$).




Figure 5.2: Summary image for the rigid polymer trial (Test 5, LEGO/ABS) highlighting the pre-test mark and the measured main damage-band width $w = 547\mu\text{m}$, comparable to the blade stack thickness and consistent with over-travel and substrate sidewall involvement. The full micrographs are provided in Appendix G (Fig. G.5).




Figure 5.3: Post-test optical micrographs of the controlled incremental rigid-sample trial (Test 6, LEGO/ABS) at $N \approx 17.7$ krPM. (a) Overview showing a continuous kerf of length $\approx 1496\mu\text{m}$ and central width $\approx 46\mu\text{m}$. (b,c) Higher-magnification views of the two kerf ends (boxes in (a)), showing a narrow initial kerf ($\sim 5\mu\text{m}$) that widens towards the centre; cutting debris is visible locally (e.g. in (b)). Scale bars: $100\mu\text{m}$ (a) and $20\mu\text{m}$ (b,c).

Using a representative kerf width of $\approx 45\mu\text{m}$, the kerf depth-to-width aspect ratio is approximately 0.3. When normalised by the rear-side film thickness ($\sim 10\mu\text{m}$), the depth-to-tool-thickness ratio is approximately 1.4. These values indicate substantial trench formation rather than a superficial surface mark. The present test did not explore maximum achievable depth, since after first kerf formation was observed, the sample was not advanced further to allow edge and sample surface inspection.

Edge attribution and post-test edge condition. Only a single kerf track was observed on the sample. No distinct narrow kerf attributable to the $\sim 4\mu\text{m}$ front-side oxide thickness was identified. Under the present conditions (rigid ABS sample and limited first-contact observability), this indicates that functional kerf formation by the ultra-thin front-side rim was not demonstrated. One plausible explanation is local fracture at (or shortly after) first rigid contact, preventing sustained engagement. This interpretation is consistent with the appearance of an additional front-side damage zone with arc length of order $8000\mu\text{m}$.

Post-test edge-length measurements showed no circumferentially global shortening of either edge. The front-side protrusion length remained essentially unchanged from $97.4\mu\text{m}$ (pre-test mean) to $97.1\mu\text{m}$ (post-test mean), and the rear-side protrusion length from $105.3\mu\text{m}$ to $104.4\mu\text{m}$. The rear-side rim exhibited slight local degradation (i.e. increased edge roughness) over a limited circumferential sector, while other regions appeared unchanged (qualitative examples in Appendix G, Fig. G.7). Together, the single kerf signature, the kerf width scale (tens of micrometres), and the observed rear-side wear pattern support the conclusion that

the rear-side ($\sim 10\mu\text{m}$ thick) rim was the dominant cutter in Test 6.

Finally, the kerf width scale is consistent in order-of-magnitude with the combination of rear-side oxide thickness and expected operational axial displacement. At $N \approx 17.7\text{kRPM}$, the axial displacement envelope (WCR_{95}) is on the order of tens of micrometres (cf. Fig. 2.4(a)), which can broaden the effective contact track of a $\sim 10\mu\text{m}$ -thick feature into a kerf width of order $40\mu\text{m}$ to $50\mu\text{m}$ under intermittent contact.

5.4. DISCUSSION AND IMPLICATIONS

Challenge 4 aimed to demonstrate integrated operation of a released blade at the minimum cutting-edge speed and to obtain a first indication of cutting engagement. The chapter provides outcomes with increasing diagnostic value. First, the released rim geometry survived rotation at and above the target cutting-edge speed without macroscopic breakaway (Test 1), establishing that the released SiO_2 structures are mechanically robust enough for high-speed operation in the developed spindle setup. Second, thread tests (Tests 2–4) demonstrated qualitative cutting engagement under compliant contact conditions while preserving edge integrity sufficiently for post-test inspection.

The rigid-sample experiments convert the chapter from a qualitative demonstration into a precision-cutting-relevant capability and limitation statement. Test 5 shows that, when first contact cannot be detected reliably (motion blur and inability to observe edge and surface simultaneously), the approach can transition abruptly into over-travel with substrate sidewall involvement and severe rim loss. In contrast, Test 6 demonstrates that controlled rigid-sample kerf formation is achievable with the present platform, since a continuous kerf of length $\sim 1.5\text{mm}$ was produced on ABS at $N \approx 17.7\text{kRPM}$, with central width $\sim 46\mu\text{m}$ and maximum trench depth $\sim 12\mu\text{m}$ to $15\mu\text{m}$, corresponding to $\text{AR}_{dlw} \approx 0.31$ and a tool-based proxy $\text{AR}_{dlt} \approx 1.39$. These values are substantial in view of the fragile rim geometry and the absence of active force control. However, the present result should not be interpreted as the maximum achievable depth, since no further post-contact advance was attempted after first kerf formation was observed.

The same rigid-sample results also sharpen rim attribution. Only one kerf track was observed, and no $\sim 4\mu\text{m}$ -scale kerf attributable to the front-side rim was identified. Under the present rigid-sample conditions, this indicates that functional kerf formation by the ultra-thin front-side rim was not demonstrated, plausibly because it fractures at first rigid contact and therefore cannot sustain engagement. Conversely, the observed kerf geometry and the post-test rear-side wear pattern support the conclusion that the rear-side ($\sim 10\mu\text{m}$ thick) rim is currently the dominant cutter. This provides a defensible basis for concluding that the present platform can perform controlled rigid cutting, while simultaneously identifying the key barrier to the long-term objective of ultra-thin cutting with edges in the sub-micrometre regime.

Two immediate implications follow for future work. First, the dominant practical limitation is observability and contact detection. Without phase-locked imaging or an independent contact cue, the approach must rely on subjective visual inference, which is inadequate for rigid-sample tests. Phase-locked strobed illumination driven from the tachometer signal, with adjustable phase and short pulse width, remains the most direct route to localising the rim in the camera view. In parallel, a compact contact indicator on the sample side (e.g. a small accelerometer on the sample holder) could provide an objective engagement cue to halt the approach at first contact and prescribe controlled post-contact advance. Second, to prevent handling-induced damage by incidental contact between the edges and the microscope stage during placement (i.e. imperfectly flat seating causes a local edge–stage touch-down), a central pedestal should be implemented that supports the blade at its hub while preventing rim contact with the microscope stage. To enable front-side and rear-side inspection without touching the rims, the pedestal should incorporate a simple clamping mechanism (e.g. a twist-and-lock feature) that allows the blade to be flipped while supported. The support structure should remain thin on both sides to avoid obstructing microscope optics and illumination.

In summary, Challenge 4 establishes an integrated proof-of-concept with a clear hierarchy of evidence: released rims survive high-speed operation; cutting engagement is achievable on compliant samples; and, critically, controlled kerf formation on a rigid polymer has been demonstrated with quantified kerf depth. The next critical step is to translate this capability to reliable, attributable cutting with the ultra-thin front-side rim by improving first-contact detectability, reducing handling-induced damage, and implementing protocols that safely explore deeper engagement without over-travel.

6

CONCLUSION AND OUTLOOK

This thesis establishes a proof-of-concept for a wafer-based, thin-film-defined circular cutting edge: a mechanically defined rim fabricated on a wafer can be rotated to cutting-edge speed and can remove material under controlled rigid contact. Enabled by a repeatable fabrication route and a characterised high-speed spindle platform, released SiO₂ rims survived operation at and above the target edge speed. Under incremental approach, a continuous kerf was produced in a rigid ABS sample at $N \approx 17.7$ krPM, with quantified trench depth 11.72 μ m to 14.71 μ m and central kerf width $\sim 46\mu$ m, corresponding to $AR_{d/w} \approx 0.31$. This result indicates substantial material removal rather than superficial marking and demonstrates that controlled rigid-sample kerf formation is achievable with the present platform.

The present work also establishes the principal limitations that must be addressed to progress from thin-film rigid cutting towards attributable ultra-thin cutting. Kerf morphology and post-test edge inspection indicate that the rear-side ($\sim 10\mu$ m) rim is currently the dominant cutter. Sustained rigid engagement of the ultra-thin front-side ($\sim 4\mu$ m) rim remains to be demonstrated under controlled loading. In addition, robust release of circumferentially continuous ultra-thin a-SiC edges was not achieved within the present process window. These outcomes position the concept as viable for rigid cutting with wafer-defined thin-film rims at cutting-edge speed, while identifying attributable ultra-thin cutting as a clearly bounded next objective.

Across the four challenges, the main outcomes of this work are:

- A practical operating envelope was established in which the spindle reaches the minimum cutting-edge speed while maintaining predominantly rotation-synchronous operational displacement in the tens of micrometres. Under the tested conditions, the displacement envelope remained below the targeted overhang length range ($L = 60\mu$ m to 120μ m), enabling survival and engagement trials without immediate centrifugal failure.
- A repeatable fabrication route was developed to produce annular blade blanks in 525 μ m-thick silicon, with LS-Precess identified as the enabling laser-cutting intervention and an explicit throughput-reliability trade-off captured in the selected recipe.
- A lid-base sacrificial masking concept was demonstrated to provide strong geometric selectivity during isotropic SF₆ etching, enabling controlled release of SiO₂ rims while preserving hub-bore integrity when the lid remains intact. For SiO₂ development blades, released overhangs in the targeted range were achieved. In contrast, circumferentially continuous ultra-thin a-SiC edges were not reliably formed.
- An integrated cutting demonstration was achieved using released SiO₂ rims. The blade survived operation at and above the target edge speed, severed a compliant sample, and produced a continuous and quantified kerf in a rigid ABS sample under incremental approach. Under these conditions, rigid cutting was attributable to the rear-side rim, while sustained rigid engagement of the ultra-thin front-side rim remains to be demonstrated.

6.1. CHALLENGE-WISE CONCLUSIONS

- **Challenge 1 (rotation to cutting-edge speed with acceptable run-out):** The developed setup reaches the minimum cutting-edge speed requirement while maintaining predominantly rotation-synchronous operational displacement at the blade edge. A practical sensitivity threshold was observed beyond ~ 18 kRPM, and quasi-static TIR was strongly assembly dependent, highlighting geometric repeatability of the clamp–shaft–bearing stack as a dominant lever for further improvement. Motor-drive excitation cannot be excluded with the present sensorless ESC, motivating future A/B comparisons with upgraded drive control.
- **Challenge 2 (femtosecond-laser fabrication of annular blanks):** Blade blanks suitable for subsequent release and high-speed trials can be produced repeatably, provided that the programmed geometry is tuned empirically to account for kerf-related dimensional bias and that handling constraints are respected. The observed laser-affected zone at the cut edge motivates a minimum underetch length on the order of $60\mu\text{m}$ to support continuous overhang formation. The current vacuum-tweezer mounting approach sets a practical lower bound on bore clearance, suggesting that a dedicated handling/pressing fixture is a direct route to improved mounting repeatability and reduced eccentricity.
- **Challenge 3 (selective underetch with hub-bore protection):** The lid–base fixture provides robust top-surface shielding, evidenced by a sharply delineated protected–exposed boundary and preservation of hub-bore geometry within present metrology when the lid remains intact. For SiO_2 process-development blades, released overhangs in the targeted range were achieved. In contrast, reliable formation of long, circumferentially continuous ultra-thin a-SiC edges was not achieved with the current process configuration, indicating that the selective process window for a-SiC must be widened, most plausibly through improved thermal management.
- **Challenge 4 (integrated cutting demonstration):** Released SiO_2 rims survived high-speed operation and demonstrated both compliant engagement (thread severing) and quantified rigid kerf formation in ABS. A first rigid trial showed that limited first-contact detectability can rapidly lead to over-travel, substrate sidewall involvement, and extensive rim loss. A subsequent trial demonstrated that controlled rigid cutting is achievable with the present platform. Under these conditions, kerf geometry and post-test wear patterns support the conclusion that the rear-side rim is currently the dominant cutter, while sustained rigid engagement of the ultra-thin front-side rim remains to be demonstrated.

6

6.2. OUTLOOK

Future work should prioritise translating the demonstrated rigid-cutting capability towards reliable, attributable ultra-thin cutting and improved process robustness. In doing so, it is useful to close the loop to state-of-the-art micro-machining: specialised dicing and micro-slotting platforms can reach depth-to-width aspect ratios $AR_{d/w}$ on the order of 4 to 20 for $\sim 10\mu\text{m}$ to $50\mu\text{m}$ kerf widths, enabled by highly optimised spindles, force/displacement control, and dedicated fixturing. The present platform achieved $AR_{d/w} \approx 0.31$ without such infrastructure and without evidence of a fundamental depth limit imposed by the $\sim 10\mu\text{m}$ rim thickness. This comparison indicates that improved first-contact control, displacement reduction, and systematic depth exploration are the primary levers for extending trench depth and narrowing kerf width in future iterations.

6.2.1. FIRST-CONTACT DETECTION, OPTICAL ACCESS, AND CONTROLLED DEPTH EXPLORATION

Rigid-sample testing revealed that visual clearance estimation under motion blur is insufficient to reliably detect first contact. In addition, as the rigid sample approaches the rim in the present configuration, the camera view of the edge–surface interface becomes progressively obstructed, reducing the ability to observe the onset of engagement precisely at the moment it occurs. The over-travel observed in the initial rigid trial demonstrates that additional advance can transition from first engagement to substrate sidewall involvement when first contact is not detected.

Future work should therefore address both synchronised visualisation and geometric optical access. Implementation of phase-locked strobed illumination driven by the tachometer signal would improve rim localisation during rotation, while a revised sample-holder geometry or alternative viewing angle could preserve a clear line-of-sight to the contact zone throughout the approach. In parallel, integration of an independent sample-side contact cue (e.g. a compact accelerometer or force-sensitive element) would enable objective

detection of first engagement. Together, these measures would permit systematic exploration of post-contact advance and maximum achievable trench depth without uncontrolled penetration.

6.2.2. ULTRA-THIN RIM ROBUSTNESS, ATTRIBUTION, AND GRADUAL FAILURE CONCEPTS

The absence of a distinct $\sim 4\mu\text{m}$ -scale kerf and the observed front-side damage zone indicate that the ultra-thin rim does not currently sustain rigid contact. Future work should therefore address the mechanical robustness and failure mode of the ultra-thin edge, while keeping loading and first-contact conditions sufficiently controlled to distinguish intrinsic fragility from externally imposed impact and bending.

A useful analogue can be found in conventional diamond dicing blades, which consist of a bonding matrix containing hard abrasive particles. Cutting proceeds by penetration of these particles into the workpiece, while progressive wear exposes new sharp particles as others fracture or are pulled out of the bond [26], [55], [56]. In contrast, a continuous ultra-thin rim has no intrinsic self-renewal capacity and is susceptible to large, uncontrolled chip-out events that locally render extended edge segments ineffective.

A longer-term research direction is therefore to explore whether the failure mode of thin-film rims can be made more gradual. Deliberate edge patterning (e.g. engineered weak links, segmentation, or controlled thickness modulation) could promote localised, contained breakaway rather than extended fracture, thereby preserving functional cutting regions after partial damage. Improved first-contact detectability and repeatable loading will be essential to unambiguously attribute cutting action to such modified ultra-thin rims.

6.2.3. SPINDLE REPEATABILITY, OPERATIONAL DISPLACEMENT, AND MOTOR-DRIVE EFFECTS

Although the operational displacement envelope is compatible with survival and initial cutting demonstrations for $L = 60\mu\text{m}$ to $120\mu\text{m}$, further reduction of operational displacement would directly influence cutting mechanics. A narrower displacement envelope would reduce the effective kerf broadening caused by lateral sweep, enabling smaller kerf widths for a given rim thickness. In addition, reduced axial and radial motion would lower bending moments and radial force components acting on the thin rims during contact.

Such improvements are not merely incremental refinements but are critical for interpretation. If ultra-thin rims fail under rigid contact, it must be possible to distinguish between intrinsic geometric fragility and externally imposed loading due to setup-induced displacement. Reducing operational displacement and clarifying motor-drive-induced excitation (e.g. via comparison with a lower-torque-ripple control strategy) would therefore enable more definitive conclusions regarding the true cutting capability and failure threshold of ultra-thin edges. The most direct routes are (i) a controlled axial/radial clearance reduction and preload strategy, (ii) characterisation and mitigation of motor-drive-induced excitation, and (iii) improved geometric repeatability of the blade clamping stack. In particular, replacing the present sensorless ESC by a drive with lower torque ripple (or an explicit A/B comparison against field-oriented control) would clarify the extent to which commutation artefacts contribute to the observed high-speed operational displacement.

6.2.4. PLASMA-ETCH PROCESS WINDOW AND TRANSLATION TO A-SiC

The lid-base concept provides a viable foundation for selective underetching, but reliable formation of circumferentially continuous ultra-thin a-SiC rims remains unresolved. Priority should be given to widening the selective process window, plausibly through improved thermal management in the fixture stack and revised etch conditions that reduce film degradation while maintaining adequate silicon undercut rates.

6.2.5. BROADER SCOPE: WAFER-DEFINED MECHANICAL EDGES AS A FABRICATION PLATFORM

Beyond the immediate next steps, the present results suggest a broader research direction: wafer processing can be used to define and replicate mechanically functional cutting edges at thickness scales defined by the thicknesses of deposited films. In doing so, the work effectively miniaturises a mechanical cutting tool to the film-thickness regime, in contrast to the predominant reliance of micro- and nanofabrication on beams, plasmas, and chemical removal processes.

In this thesis, Si/SiO₂/a-SiC were selected primarily for process compatibility and availability within a semiconductor workflow. However, the underlying concept is not intrinsically limited to these materials. Subject to suitable deposition, adhesion, and release strategies, analogous rotating rims could be envisaged in other hard or functional thin films, such as carbides, diamond-like carbon, or selected metallic layers. The key requirement is the ability to fabricate, release, and mechanically support a circumferential edge with sufficient

hardness and fracture resistance for controlled engagement under rotation.

A further distinction of mechanically defined cutting is that material removal is governed primarily by contact mechanics and relative hardness rather than by chemical selectivity or sputter yield. As such, rotating thin-film edges may offer a complementary structuring route for materials or multilayer stacks where plasma etching, ion-beam milling, or wet processing face limitations related to etch chemistry, selectivity, or surface charging. The present demonstration in ABS provides an initial indication that polymer substrates can be structured under controlled rotation, although systematic exploration of additional material classes remains future work.

While attributable sub-micrometre cutting has not yet been demonstrated, this work establishes that wafer-defined rims can survive cutting-edge speed and perform quantified rigid material removal. With improved control over first contact, operational displacement, and material selection, wafer-defined rotating edges may constitute a complementary manufacturing route for continuous, mechanically defined structuring under rotation alongside established beam- and plasma-based techniques.

A

SIGNAL PROCESSING AND ROBUSTNESS CHECKS

This appendix provides the signal-processing details underlying the operational displacement metrics (RSR and WCR) reported in Section 2.6. The key challenge is that the LDV output is a *velocity* signal measured during rotation, while the spindle speed can drift slightly within a 7 s record due to the ESC. To obtain drift-robust displacement estimates without time-integration drift, velocity is converted to displacement in the *order domain* using the tachometer signal as an angular reference.

A.1. ORDER-DOMAIN DISPLACEMENT RECONSTRUCTION FROM LDV VELOCITY

The LDV measures the surface velocity $v(t)$ along the measurement axis. Using the tachometer, each record is segmented into individual revolutions and re-sampled onto a uniform angular grid $\theta \in [0, 2\pi)$. For a given revolution, the velocity can be expressed as a truncated Fourier series in rotation orders:

$$v(\theta) \approx \sum_{m=1}^M V_m e^{jm\theta}, \quad (\text{A.1})$$

where m denotes the order and V_m is the complex velocity coefficient at order m .

Displacement is obtained *in the order domain* via harmonic-by-harmonic division by $jm\omega$, avoiding time-domain integration drift:

$$X_m = \frac{V_m}{jm\omega}, \quad \omega = 2\pi f_{\text{rot}}, \quad (\text{A.2})$$

where f_{rot} is the rotational frequency associated with the record and X_m is the complex displacement coefficient.

The reconstructed displacement waveform for a truncation order M is then

$$x_M(\theta) = \sum_{m=1}^M \Re \{ X_m e^{jm\theta} \}. \quad (\text{A.3})$$

The cumulative peak-to-peak displacement used in the convergence analysis is computed directly from the reconstructed waveform:

$$pp(M) = \max_{\theta} (x_M(\theta)) - \min_{\theta} (x_M(\theta)). \quad (\text{A.4})$$

A.2. OPERATING POINT DEFINITION UNDER RPM DRIFT

Because the ESC-controlled speed can vary slightly within a record, each operating point is represented by the *median RPM* of that record (as used in Section 2.6). The within-record speed variation is summarised by the 10–90% RPM span, ΔRPM_{10-90} , computed from the instantaneous RPM trace derived from the tachometer.

For the axial blade-radii dataset (Blade ON, $r = 15.5$ mm to 17.5 mm, $n = 29$ records), the median span was $\Delta\text{RPM}_{10-90} \approx 116$ RPM (median $\approx 1.27\%$ of the record median RPM), with a worst-case span of 458 RPM (3.23%). For the radial rim dataset ($n = 8$ records), the median span was ≈ 146 RPM (1.39%) with a worst-case span of 564 RPM (2.71%). A drift-rate metric $d\text{RPM}/dt$ (computed in the KPI pipeline) was typically on the order of 5 RPM/s–8 RPM/s (axial median 5.31 RPM/s, radial median 7.78 RPM/s), with occasional larger excursions (axial maximum 42.0 RPM/s, radial maximum 30.3 RPM/s). These magnitudes motivate (i) plotting KPIs against the median RPM per record and (ii) using order-domain reconstruction, which remains well-defined under modest speed drift.

A.3. BOOTSTRAP PROCEDURE FOR UNCERTAINTY ESTIMATES

Uncertainty bands for peak-to-peak metrics are estimated by bootstrap resampling over revolutions within each record. Let R be the number of available revolutions in a record. For each bootstrap replicate $b = 1, \dots, B$ (with $B = 1000$ in this work), a new set of R revolutions is drawn *with replacement* from the original revolutions, and the corresponding displacement coefficients $\{X_m\}$ and derived metrics are re-computed. For any scalar metric y (e.g., $pp(M)$, RSR p-p, WCR_{95}), the 95% confidence interval is taken from the empirical bootstrap distribution as

$$y_{\text{lo}95} = \text{prctile}(y^{(b)}, 2.5), \quad y_{\text{hi}95} = \text{prctile}(y^{(b)}, 97.5). \quad (\text{A.5})$$

In the main text, confidence intervals are not used as a substitute for run-to-run repeatability (only one record per operating point is available); rather, they quantify the within-record variability of the reconstructed KPI under revolution-level resampling.

A.4. SELECTION OF RECONSTRUCTION ORDER M_{recon}

The peak-to-peak displacement $pp(M)$ depends on the number of included orders. To determine a suitable reconstruction order, a convergence analysis is performed by evaluating $pp(M)$ for $M = 1, \dots, M_{\text{max}}$ (with $M_{\text{max}} \leq 20$). The incremental relative change between successive truncation levels is defined as

$$\text{inc}(M) = \frac{|pp(M) - pp(M-1)|}{pp(M)}, \quad M \geq 2. \quad (\text{A.6})$$

A stability threshold of 2% sustained over three consecutive steps is used:

$$\text{inc}(M), \text{inc}(M+1), \text{inc}(M+2) < 0.02. \quad (\text{A.7})$$

The selected reconstruction order M_{recon} is taken as the *smallest* M satisfying (A.7). This choice ensures that the reported peak-to-peak metrics are insensitive to including additional orders beyond M_{recon} .

Fig. A.1 reports the resulting M_{recon} values across RPM and blade radii. For the blade-radii measurements, M_{recon} was consistently ≥ 5 and typically 5–6 across the explored RPM range, indicating that several harmonics are required to obtain stable peak-to-peak displacement estimates. (No cases were observed where the convergence criterion was unmet within the evaluated order range.)

A.5. HARMONIC DISTORTION RATIO (HDR) AND INTERPRETATION

To summarise the relative contribution of higher orders beyond $1\times$, a harmonic distortion ratio is defined as

$$\text{HDR}_{2\dots M_{\text{recon}}} = \frac{\sum_{m=2}^{M_{\text{recon}}} |X_m|}{\max(\epsilon, |X_1|)}, \quad (\text{A.8})$$

where ϵ is a small numerical constant to avoid division by zero. This metric increases when higher-order content becomes more prominent relative to the fundamental ($1\times$) component.

Fig. A.2 shows $\text{HDR}_{2\dots M_{\text{recon}}}$ versus RPM for the axial blade-radii dataset. Across the overlapping RPM range where multiple radii are available, $\text{HDR}_{2\dots M_{\text{recon}}}$ decreases with RPM, indicating that the operational motion becomes increasingly dominated by the $1\times$ component at higher speed. A tendency for higher HDR values closer to the rim is observed within the overlapping range; at the highest RPM only rim data are available, and cross-radius comparison is therefore not possible there.




Figure A.1: M_{recon} selected by the convergence criterion of (A.7) as a function of RPM for the axial blade-radii dataset. Each point corresponds to one 7 s record plotted at its median RPM; n denotes the number of records per radius.




Figure A.2: Harmonic distortion ratio $HDR_{2...M_{recon}}$ (A.8) versus RPM for the axial blade-radii dataset. Each point corresponds to one 7 s record plotted at its median RPM; n denotes the number of records per radius.

B

SUPPLEMENTARY RESULTS FOR CHALLENGE

1

This appendix provides supplementary operational-displacement results that support, but are not required for, the main narrative in Challenge 1. The main text focuses on blade-rim behaviour. Here, baseline measurements at the shaft axis ($r = 0$ mm) and the full set of axial blade-radii results are reported for completeness. As in Section 2.6, each marker corresponds to one 7 s record and is plotted at the measured median RPM of that record; n denotes the number of records in each series.

B.1. BASELINE OPERATIONAL DISPLACEMENT AT THE SHAFT AXIS ($r = 0$ mm)

Fig. B.1 compares the operational displacement at the shaft axis for three baseline configurations (blade removed): pretension and clamping attachments were varied to assess their influence on motion near the rotation axis. Across the tested speed range, both the repeatable rotation-synchronous component (RSR) and the per-revolution severity measure (WCR_{95}) reduce when pretension and clamping attachments are applied. The separation between WCR_{95} (dashed) and RSR (solid) remains modest, indicating that the axis motion is predominantly rotation-synchronous within these baseline tests.

B.2. FULL AXIAL BLADE-RADII RESULTS (RSR AND WCR_{95})

Fig. B.2 shows the full set of axial blade-radii results for $r = 15.5$ mm to 17.5 mm. Panel (a) reports RSR p–p, and panel (b) reports WCR_{95} p–p. In the overlapping RPM range (where multiple radii are available), operational displacement generally increases towards the rim. The vertical reference line indicates the minimum cutting-edge-speed requirement (15.9 kRPM, corresponding to 30 m/s at $r = 18$ mm). At higher speeds, only the rim series extends to the maximum tested RPM.

B.2.1. ORDER-DOMAIN DECOMPOSITION ACROSS BLADE RADII

Fig. B.3 and Fig. B.4 complement the main-text RSR/ WCR results by showing how the operational displacement is distributed across integer orders. They are referenced in Section 2.7 to support the observed low-order dominance and the transition from a pronounced $2\times$ component at lower/mid RPM to increasing $1\times$ content near the high-speed upturn.

B

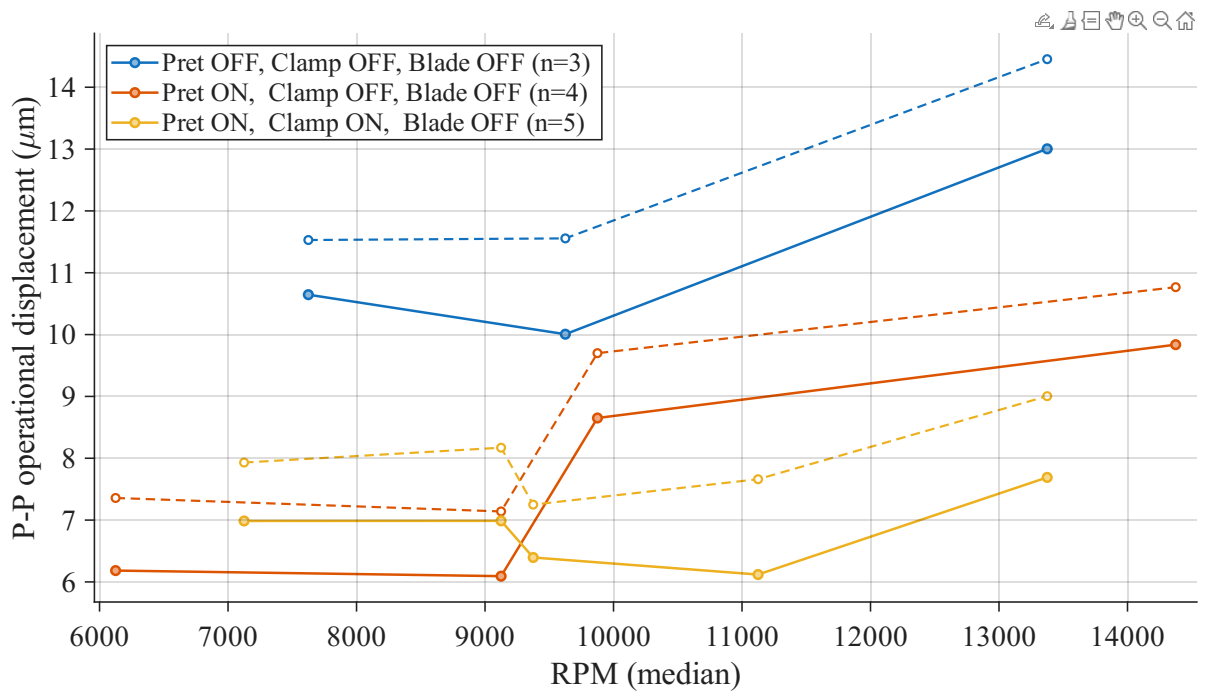


Figure B.1: Operational displacement at the shaft axis ($r = 0$ mm) for baseline configurations without the blade. Solid curves denote RSR; dashed curves denote WCR_{95} (95th percentile of per-revolution peak-to-peak values within a record). Each point corresponds to one 7 s record plotted at its median RPM; n indicates the number of records per configuration.

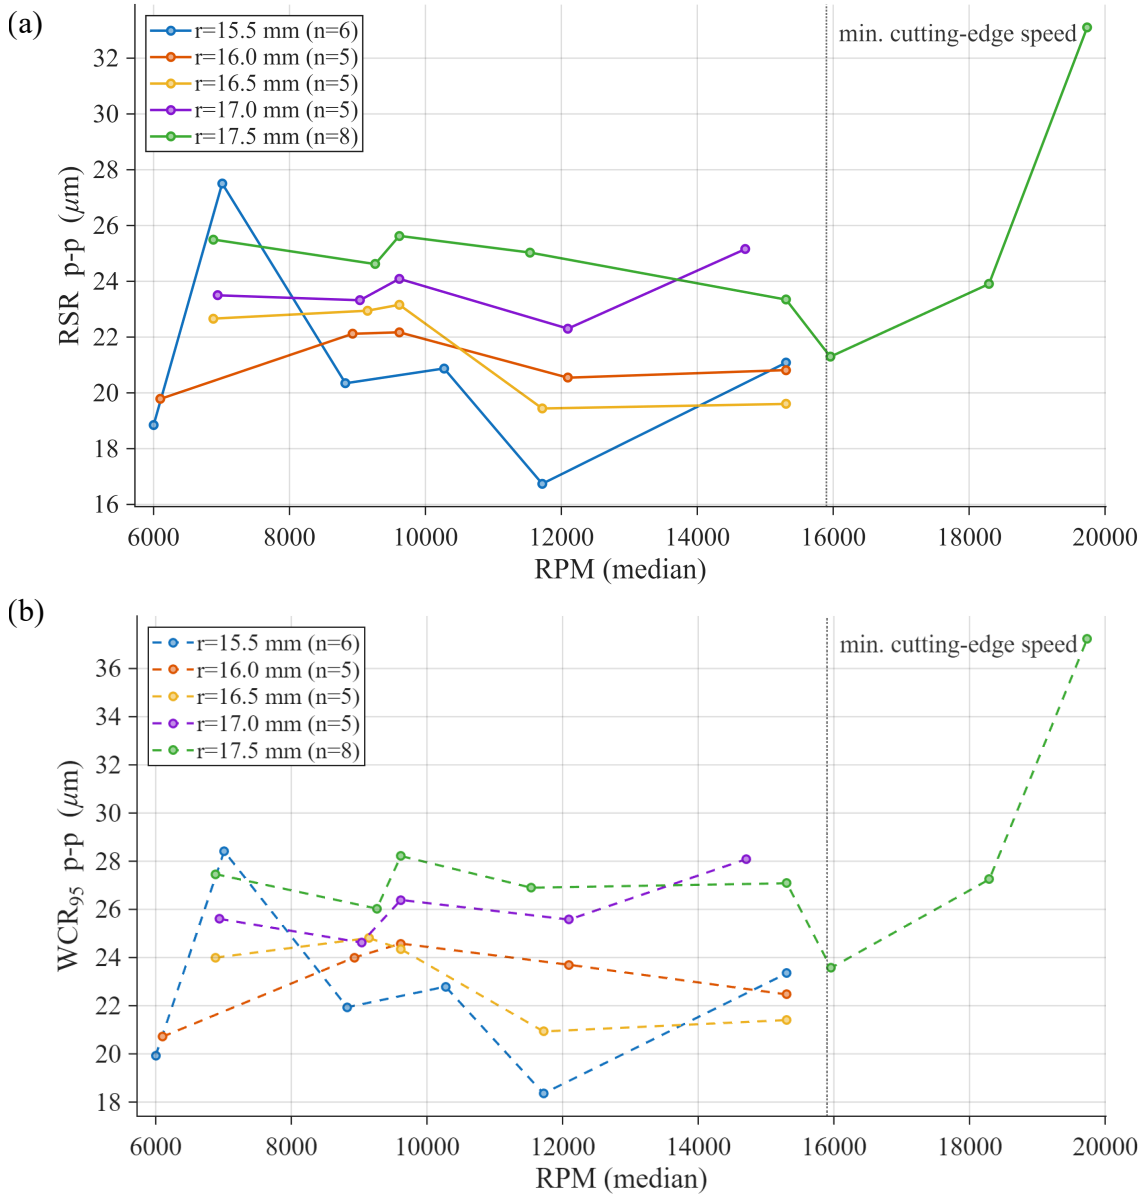


Figure B.2: Full axial blade-radii operational displacement results. (a) RSR p-p versus rotational speed for $r = 15.5$ mm to 17.5 mm. (b) WCR₉₅ p-p versus rotational speed for the same radii. Each point corresponds to one 7 s record plotted at its median RPM; n indicates the number of records per radius. The vertical reference line indicates the minimum cutting-edge-speed requirement (15.9 kRPM, corresponding to 30 m/s at $r = 18$ mm).

B




Figure B.3: Fractional amplitude contribution of integer order components versus rotational speed for axial LDV measurements at multiple blade radii ($r = 16.0$ mm, 16.5 mm, 17.0 mm, 17.5 mm). For each 7 s record, the order-tracked displacement spectrum was decomposed into integer orders $k = 1 \dots M$, and the plotted quantity is the normalised contribution $|X_k|/\sum_{i=1}^M |X_i|$ (computed over the same order range). Each marker corresponds to one record plotted at its measured median RPM; marker shape encodes radius.

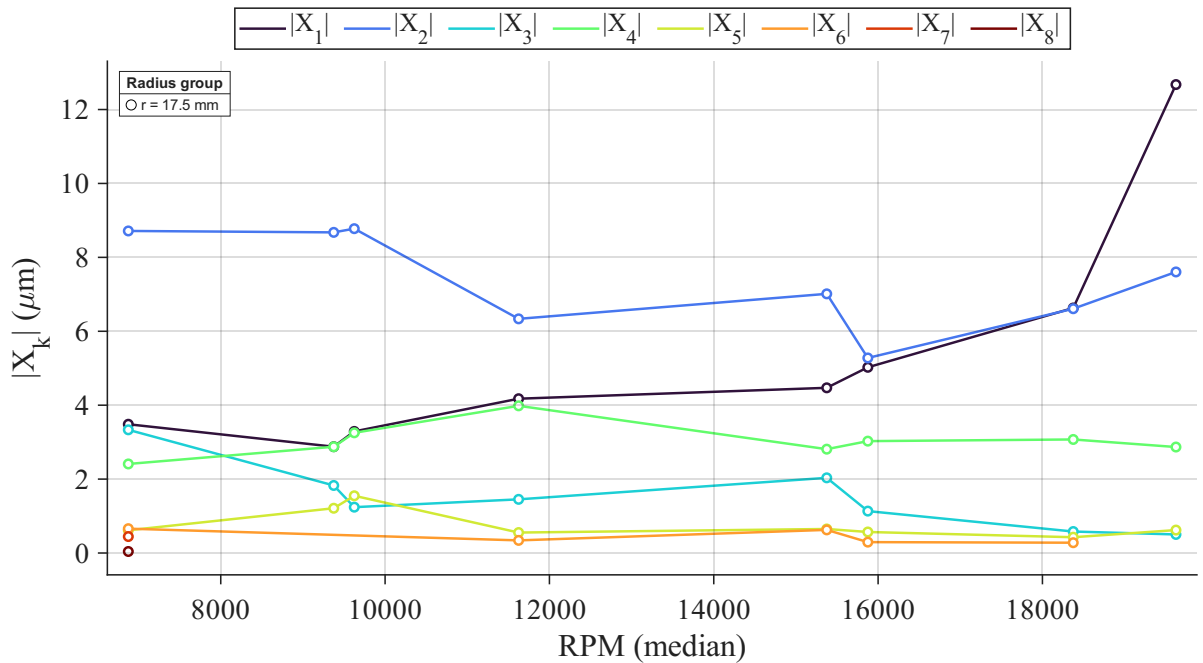


Figure B.4: Per-order displacement amplitude versus rotational speed at the blade rim ($r = 17.5$ mm, axial LDV). Curves show the absolute amplitudes $|X_k|$ of integer orders $k = 1 \dots M$ (with M chosen per record according to the convergence criterion described in Section A.4). Each marker corresponds to one 7 s record plotted at its measured median RPM.

C

SUPPLEMENTARY FIGURES FOR CHALLENGE

2

C.1. QUALITATIVE INFLUENCE OF LASER POWER ON NEAR-EDGE MORPHOLOGY IN LS-PRECESS MODE

During recipe refinement (Chapter 3), laser power was iterated within the policy-limited range (60 % to 80 %) while keeping the remaining scan settings constant. The aim was not a metrological optimisation, but a practical selection of settings that (i) reliably achieve through-cuts and (ii) minimise pronounced near-edge irregularities that could locally interrupt overhang formation in the subsequent underetch step (Challenge 3).

Fig. C.1 shows representative top-surface kerf segments. Differences across the tested range are modest and vary locally. However, within the inspected samples the 80 % condition tended to exhibit a straighter perimeter with fewer pronounced local notches than 60 % and 70 %. On this qualitative basis, 80 % was adopted for the final production recipe reported in Chapter 3.




Figure C.1: Qualitative effect of laser power on the top-edge morphology in LS-Precess mode. Keyence VHX-6000 micrographs of representative kerf segments cut at constant scan speed (200 mm s^{-1}), outlines (two), and outline pitch (0.01 mm). (a) 60 %, (b) 70 %, (c) 80 %. Red curves indicate manually traced kerf edges. Scale bars: $20\mu\text{m}$.

C.2. MANUAL MOUNTING LIMITATION WHEN USING A VACUUM TWEEZER

Bore-fit tuning in Chapter 3 was constrained by the mounting method. Blades were handled and mounted using a vacuum tweezer to avoid contacting the fragile outer edge. This approach provides a compliant, single-point grip via a soft suction cup. While protective, it limits controllable alignment and the axial seating force that can be applied during assembly.

As the bore clearance decreases, initial contact between the bore and the shaft seat can generate frictional moments that tilt the disk. Because the grip point is compliant and offset from the bore contact, correcting tilt and applying additional seating force becomes increasingly difficult without touching the perimeter. This practical constraint sets the lower bound on clearance in the stepwise fit tests and motivates the proposed

future improvement of a simple mounting/pressing fixture.

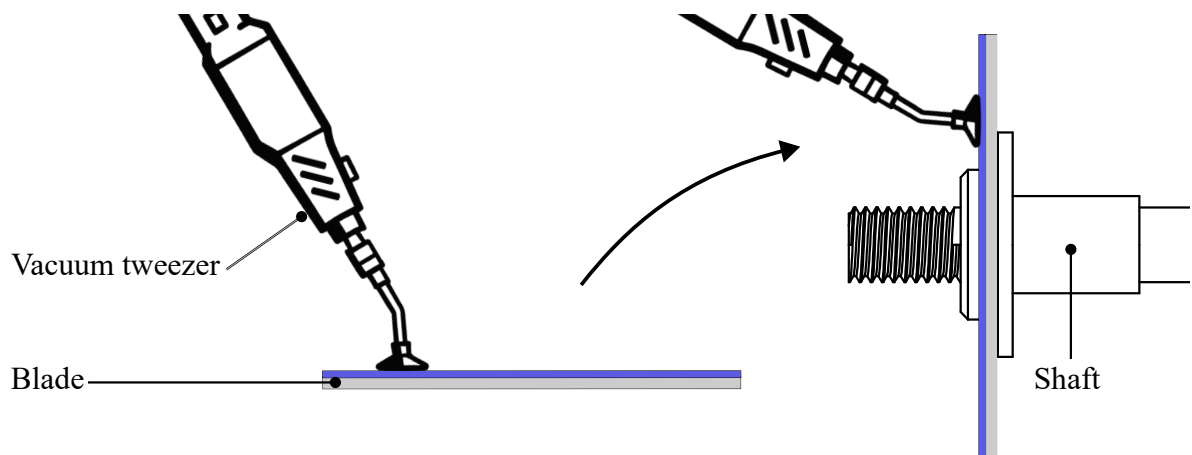


Figure C.2: Schematic of blade mounting using a vacuum tweezer. The suction cup avoids contact with the outer edge but provides a compliant, single-point grip, which limits controllable alignment and axial assembly force as the bore engages the shaft seat.

D

INITIAL BENDING AND ROTATIONAL STIFFENING

This appendix documents the modelling assumptions and parameter studies used to estimate (i) the initial out-of-plane deflection of the released cutting edge due to residual stress gradients, and (ii) the reduction of this deflection under rotation (centrifugal stiffening). The cutting edge is modelled locally as a cantilever-like beam with rectangular cross-section, length L , thickness d (film thickness) and in-plane width w . A one-mode Rayleigh–Ritz (energy) formulation is used for the rotating-beam estimate. The full derivation is provided separately in Appendix E, while this appendix summarises the governing relations and their numerical evaluation. All numerical evaluations presented in this appendix were carried out in MATLAB. The script used to generate Tables D.1–D.4 is provided in Appendix F.

D.1. THEORETICAL BACKGROUND: INITIAL DEFLECTION

A thin film deposited on a substrate generally exhibits residual stress arising from two main contributions, an *extrinsic* component commonly associated with thermal-mismatch or externally imposed effects, and an *intrinsic* component originating from the film growth process itself [44]. The extrinsic part is typically uniform over the film thickness and is dominated by thermal-mismatch stress between the film and substrate.

For a film bonded to a substrate, a mismatch in thermal expansion coefficients gives rise to a nearly uniform in-plane stress that is well described by classical thermal-mismatch theory [44]. This extrinsic contribution primarily sets the average residual stress level in the film. Intrinsic stress, by contrast, depends on the film microstructure and deposition conditions (deposition rate, temperature, pressure, incorporation of impurities, grain structure, etc.) and is often non-uniform through the thickness [44]. It is therefore convenient to




Figure D.1: Schematic diagram of residual stress in a thin film of thickness d , decomposed into a uniform component σ_0 and a linear gradient component σ_1 (adapted from [53]).




Figure D.2: Curvature of a free-standing thin film of thickness d produced by a stress gradient. The sketch shows the approximate stress profile and the corresponding relaxed strain ϵ_x across the film thickness (adapted from [44])

express the total residual stress in the film as a polynomial expansion in the thickness direction y [52]:

$$\sigma_f(y) = \sum_{k=0}^{\infty} \sigma_k \left(\frac{y}{d/2} \right)^k \approx \sigma_0 + \sigma_1 \left(\frac{y}{d/2} \right), \quad (\text{D.1})$$

where d is the film thickness, σ_0 is the average (uniform) stress and σ_1 represents the first-order stress-gradient term, higher-order terms are neglected. As illustrated schematically in Fig. D.1, the total residual stress σ_f can thus be viewed as the superposition of a uniform component σ_0 and a linearly varying component σ_1 across the film thickness. The average stress is strongly influenced by extrinsic effects, whereas the stress gradient is mainly governed by intrinsic mechanisms.

When a stressed film remains bonded to a thick substrate and the film–substrate system is free to move, the residual stress is replaced by a relaxed strain that bends the film and substrate and creates a curvature with radius R . This curvature $1/R$ can be related to the total residual film stress via the Stoney equation. The Stoney equation provides the *average* residual stress in the unrelaxed film directly from the measured wafer curvature and the known substrate properties, without requiring any knowledge of the film's properties. This approach is known as the substrate curvature (wafer bending) method [44], [57], [58].

However, the stress obtained from the Stoney equation is an *average* value, since it results from integrating the stress over the thickness of the film. It does not reveal whether a stress gradient is present, nor its magnitude. In this work, the ultra-thin cutting edge is released from the Si substrate by under-etching, forming a free-standing cantilever-like overhang. For such structures, the out-of-plane curvature is governed primarily by the *stress gradient* rather than the average stress [50]. Uniform stress may contribute to local boundary rotation and near-anchor deformation. This contribution however is small and difficult to measure accurately. Therefore, it is often neglected in released cantilever analysis [44], [52], [59].

The mechanism, which bends free-standing cantilevers, can be understood by considering a monolayer film that is initially clamped to the substrate. Fig. D.2 depicts such a thin-film cantilever with thickness d , clamped to a substrate. Before release, the film carries both an average stress and a stress gradient but is constrained, so there is no in-plane strain ($\epsilon_x = 0$). If the film were released while artificially constrained to remain flat, the average stress would relax into a uniform in-plane strain, but no curvature would develop. At this point the beam still contains the stress gradient, in the Fig. D.2 compression at the bottom and tension at the top. Once the flatness constraint is removed, the remaining stress gradient relaxes into a strain gradient, causing the beam to bend towards the tensile side into a curved configuration. The resulting curvature therefore directly reflects the stress gradient in the film. In this project, a stress gradient in the film could thus result in an out-of-plane deflection of the ultra-thin cutting edge.

Table D.1: Calculated initial tip deflection h_0 (in nm) as a function of cantilever length L and stress gradient $\Delta\sigma/d$, obtained from (D.5) for a-SiC with $E = 210$ GPa.

| L (μm) | Stress gradient ($\text{MPa}\mu\text{m}^{-1}$) | | | | | |
|---------------------|--|---------|---------|----------|----------|----------|
| | 50 | 100 | 150 | 200 | 250 | 300 |
| 5 | 3 nm | 6 nm | 9 nm | 12 nm | 15 nm | 18 nm |
| 10 | 12 nm | 24 nm | 36 nm | 48 nm | 60 nm | 71 nm |
| 40 | 191 nm | 381 nm | 571 nm | 762 nm | 952 nm | 1143 nm |
| 60 | 429 nm | 857 nm | 1286 nm | 1714 nm | 2143 nm | 2571 nm |
| 80 | 762 nm | 1524 nm | 2286 nm | 3048 nm | 3810 nm | 4571 nm |
| 100 | 1191 nm | 2381 nm | 3571 nm | 4762 nm | 5952 nm | 7143 nm |
| 120 | 1714 nm | 3429 nm | 5143 nm | 6857 nm | 8571 nm | 10286 nm |
| 150 | 2679 nm | 5357 nm | 8036 nm | 10714 nm | 13393 nm | 16071 nm |

D

D.2. ESTIMATING THE INITIAL TIP DEFLECTION OF THE CUTTING EDGE

In this work the cutting edge is modelled locally as a cantilever beam with rectangular cross-section, length L , thickness d and width w . This one-dimensional approximation neglects the additional stiffness provided by the surrounding annular disk. The real cutting edge behaves more like a curved plate segment, so the cantilever model is expected to overestimate the deflection and can be regarded as a conservative upper bound. The initial (non-rotating) tip deflection is denoted $h_0 \equiv h(\Omega = 0)$.

Assuming that only the stress-gradient term contributes to bending and that it varies linearly over the thickness of the beam, the stress gradient and curvature are related by [44]

$$\Delta\sigma = E \frac{d}{R}, \quad \Delta\epsilon = \frac{d}{R}, \quad (\text{D.2})$$

where R is the radius of curvature and E is the Young's modulus of the film. For a cantilever of length L with maximum out-of-plane deflection h_0 , simple geometric considerations yield the curvature

$$R = \frac{L^2}{2h_0}, \quad (\text{D.3})$$

which is strictly valid in the small-deflection regime ($h_0 \ll L$), i.e. when the deflected shape can be approximated by a shallow circular arc. Combining (D.2) and (D.3) gives the stress gradient per unit thickness as

$$\frac{\Delta\sigma}{d} = E \frac{2h_0}{L^2}. \quad (\text{D.4})$$

Rearranging gives an expression for the initial tip deflection in terms of the stress gradient,

$$h_0 = \frac{\Delta\sigma}{d} \frac{L^2}{2E}. \quad (\text{D.5})$$

The parametric studies have been carried out using stress-gradient values in the range $\Delta\sigma/d = 50 \text{ MPa}\mu\text{m}^{-1}$ to $300 \text{ MPa}\mu\text{m}^{-1}$. This range was chosen to cover the order of magnitude reported in the literature for SiC-based and similar brittle films, where stress gradients from near zero up to approximately $280 \text{ MPa}\mu\text{m}^{-1}$ have been measured depending on deposition conditions [44], [51], [53], [54]. The upper bound of $300 \text{ MPa}\mu\text{m}^{-1}$ is therefore slightly conservative.

Using (D.5), the initial tip deflection was calculated for cantilever lengths from $5 \mu\text{m}$ to $150 \mu\text{m}$ and a Young's modulus $E = 210$ GPa for amorphous SiC [60]. Representative values are shown in Table D.1. The results indicate that significant initial tip deflection only occurs for relatively long beams and high stress gradients. For beams with $L \leq 40 \mu\text{m}$ the predicted deflection remains below $1 \mu\text{m}$ for stress gradients up to $250 \text{ MPa}\mu\text{m}^{-1}$, and only just exceeds $1 \mu\text{m}$ at the highest value considered. At intermediate lengths ($L \approx 60 \mu\text{m}$ to $80 \mu\text{m}$) the deflection lies in the sub-micrometre to few-micrometre range. For $L \geq 100 \mu\text{m}$ the tip deflection increases to several micrometres at high stress gradients, reaching $\mathcal{O}(10 \mu\text{m})$ for $L = 150 \mu\text{m}$ at the upper-end stress gradients.

To express the curvature in a dimensionless form, the ratio h_0/L was also evaluated (Table D.2). Over the full parameter range this ratio remains below 1% for the shortest beams ($L \leq 10 \mu\text{m}$) and below about 3%

Table D.2: Calculated dimensionless initial tip deflection h_0/L (in %) as a function of cantilever length L and stress gradient $\Delta\sigma/d$, obtained from (D.5) for a-SiC with $E = 210$ GPa.

| L (μm) | Stress gradient ($\text{MPa}\mu\text{m}^{-1}$) | | | | | |
|---------------------|--|------|------|------|------|-------|
| | 50 | 100 | 150 | 200 | 250 | 300 |
| 5 | 0.1% | 0.1% | 0.2% | 0.2% | 0.3% | 0.4% |
| 10 | 0.1% | 0.2% | 0.4% | 0.5% | 0.6% | 0.7% |
| 40 | 0.5% | 1.0% | 1.4% | 1.9% | 2.4% | 2.9% |
| 60 | 0.7% | 1.4% | 2.1% | 2.9% | 3.6% | 4.3% |
| 80 | 0.9% | 1.9% | 2.9% | 3.8% | 4.8% | 5.7% |
| 100 | 1.2% | 2.4% | 3.6% | 4.8% | 6.0% | 7.1% |
| 120 | 1.4% | 2.9% | 4.3% | 5.7% | 7.1% | 8.6% |
| 150 | 1.8% | 3.6% | 5.4% | 7.1% | 8.9% | 10.7% |

Table D.3: Calculated dimensionless initial tip deflection h/L (in %) as a function of cantilever length L and stress gradient $\Delta\sigma/d$. Values are obtained from (D.5) for an SiO_2 film $E = 70$ GPa.

| L (μm) | Stress gradient ($\text{MPa}\mu\text{m}^{-1}$) | | | | | |
|---------------------|--|-------|-------|-------|-------|-------|
| | 50 | 100 | 150 | 200 | 250 | 300 |
| 5 | 0.2% | 0.4% | 0.5% | 0.7% | 0.9% | 1.1% |
| 10 | 0.4% | 0.7% | 1.1% | 1.4% | 1.8% | 2.1% |
| 40 | 1.4% | 2.9% | 4.3% | 5.7% | 7.1% | 8.6% |
| 60 | 2.1% | 4.3% | 6.4% | 8.6% | 10.7% | 12.9% |
| 80 | 2.9% | 5.7% | 8.6% | 11.4% | 14.3% | 17.1% |
| 100 | 3.6% | 7.1% | 10.7% | 14.3% | 17.9% | 21.4% |
| 120 | 4.3% | 8.6% | 12.9% | 17.1% | 21.4% | 25.7% |
| 150 | 5.4% | 10.7% | 16.1% | 21.4% | 26.8% | 32.1% |

for $L = 40\mu\text{m}$. For intermediate lengths ($L = 60\mu\text{m}$ to $80\mu\text{m}$) h_0/L ranges from approximately 0.7% to 4.3% and reaches at most 5.7% at $L = 80\mu\text{m}$ for the highest stress gradient considered. Only for the longest beams ($L \geq 100\mu\text{m}$) and the largest stress gradients does h_0/L exceed 5%, with a maximum value of 10.7% at $L = 150\mu\text{m}$ and $\Delta\sigma/d = 300\text{MPa}\mu\text{m}^{-1}$. Thus, only the longest beams subjected to the largest plausible stress gradients exhibit curvature on the order of 10% of their length, whereas beams with $L \lesssim 40\mu\text{m}$ deflect by at most a few percent of their length and can be regarded as nearly straight. For these latter cases the small-deflection approximation underlying (D.3) is well satisfied. For the most extreme combinations (large L and large $\Delta\sigma/d$) the estimates become approximate but remain useful as conservative upper bounds.

D.2.1. INITIAL BENDING ESTIMATES FOR THE SiO_2 EDGE (COMPARISON CASE)

Table D.3 reports the estimated dimensionless initial tip deflection h_0/L for a $4\mu\text{m}$ TEOS SiO_2 film, evaluated over the same parameter ranges used for the a-SiC case ($L = 5\mu\text{m}$ to $150\mu\text{m}$ and $\Delta\sigma/d = 50\text{MPa}\mu\text{m}^{-1}$ to $300\text{MPa}\mu\text{m}^{-1}$). The same stress-gradient range is used here for *direct comparison* with the a-SiC case, and is not intended to imply that TEOS SiO_2 necessarily exhibits this full range under the deposition conditions used for the inventory wafer.

Since $h_0 \propto 1/E$ in (D.5), using $E \approx 70$ GPa for SiO_2 [44] implies initial curvature that is approximately three times larger than for a-SiC ($E \approx 210$ GPa) at identical $\Delta\sigma/d$. Consistent with this scaling, the SiO_2 values in Table D.3 exceed the corresponding a-SiC values at all lengths and stress gradients. Notably, even at $L = 40\mu\text{m}$ the curvature can reach $h_0/L = 8.6\%$ at the upper-end stress gradients, and for $L = 150\mu\text{m}$ it can reach $h_0/L = 32.1\%$. These regimes no longer satisfy the small-deflection assumption underlying (D.3), the SiO_2 estimates should therefore be interpreted as conservative, order-of-magnitude bounds rather than precise predictions.

Finally, the actual cutting edge forms part of a continuous annular ring rather than an isolated cantilever; its stiffness is therefore higher and the true deflections are expected to be smaller than these estimates. The values in Tables D.1 and D.2, and Table D.3, should therefore be interpreted as conservative upper bounds on the initial tip deflection of the cutting edge.

D.3. ROTATIONAL STIFFENING MODEL

The initial-deflection analysis in the previous paragraph provides the static tip displacement $h_0(L)$ of the ultra-thin cutting edge, modelled as a cantilever beam, at zero rotational speed, as a function of beam length and stress gradient in the film. During operation, however, the blade is spun at high rotational speeds and the resulting centrifugal forces are expected to tension the overhanging beam and reduce its deflection.

To estimate this effect, the ultra-thin cutting edge is again idealised as a centrifugally stiffened cantilever beam. The beam has length L , in-plane width w and thickness d , and is clamped at a radius R_E from the rotation axis (see Fig. 4.5). For a blade with outer radius r_o , which is set by the femtosecond laser machine, an overhang of length L that protrudes radially from the blade body satisfies

$$r_o = R_E + L \Rightarrow R_E = r_o - L.$$

The stress gradient is represented by an initial tip deflection h_0 , approximated by the classical cubic deflection shape of a prismatic Euler-Bernoulli cantilever under a tip load. This shape function satisfies the essential clamped boundary conditions (zero deflection and slope at $x = 0$) and the zero-moment condition at the free tip ($x = L$). It is also close to the first bending mode of a cantilever and is therefore a suitable single-mode trial function for a Rayleigh–Ritz estimate. Alternative smooth shape functions (e.g. $(x/L)^2$ or the exact first eigenmode) were considered. These mainly change the numerical stiffening coefficient and do not affect the qualitative Ω^2 dependence. Given the uncertainty in the true stress gradient induced curvature, the present choice is deemed sufficiently accurate for the intended crude estimate.

D.4. NUMERICAL EVALUATION OF ROTATIONAL STIFFENING

Section D.2 provides the initial (non-rotating) tip deflection $h_0(L)$ of the released cutting edge due to a residual stress gradient. During operation, the blade is spun at high rotational speed and centrifugal loading induces axial tension in the overhang, which increases the effective transverse stiffness and reduces the tip deflection. A one-mode Rayleigh–Ritz formulation was used to obtain a closed-form estimate for this centrifugal stiffening. The full derivation and assumptions are provided in Appendix E (Rayleigh–Ritz derivation). Here, the governing relations are summarised and evaluated for the design-relevant parameter range. The rotating-beam estimate yields an effective, speed-dependent tip stiffness $k_{\text{eff}}(L, \Omega)$ and the corresponding reduced tip deflection $h(\Omega)$. Using the notation of Appendix E, the stiffening factor can be written compactly as

$$\frac{k_{\text{eff}}(L, \Omega)}{k_0(L)} = 1 + \alpha, \quad \frac{h(\Omega)}{h_0} = \frac{k_0}{k_{\text{eff}}} = \frac{1}{1 + \alpha}, \quad (\text{D.6})$$

with

$$\alpha = \frac{\rho A \Omega^2 L^3 (27L + 35R_E)}{280EI}, \quad A = wd, \quad I = \frac{wd^3}{12}, \quad R_E = r_o - L, \quad (\text{D.7})$$

where r_o is the blade outer radius (set by the femtosecond laser machine), R_E is the radial position of the clamped end of the overhang (Fig. 4.5), ρ is the film density, and E is the Young's modulus. Importantly, α depends on geometry, material properties, and rotational speed, but not on the stress gradient. Therefore, for a given L and Ω , (D.6) applies uniformly to all initial deflections h_0 reported in Section D.2.

D.4.1. EVALUATION FOR THE A-SiC DESIGN CASE

For the a-SiC design case, (D.6) and (D.7) were evaluated using the reference geometry and operating-speed assumption of Section 4.2.2: $w = 10\mu\text{m}$, $d = 96\text{nm}$, $r_o = 18\text{mm}$, and $N = 20000\text{ RPM}$ ($\Omega = 2\pi N/60$). Material properties were taken as $E = 210\text{ GPa}$, and $\rho = 2900\text{ kg/m}^3$ [60]. Table D.4 summarises the resulting stiffening factors k_{eff}/k_0 and the associated relative deflection reduction $\Delta h/h_0 = 1 - h(\Omega)/h_0 = 1 - 1/(k_{\text{eff}}/k_0)$.

Table D.4 shows that centrifugal stiffening is negligible for short overhangs ($L \leq 40\mu\text{m}$), where the predicted deflection reduction remains below 2% at 20000 RPM. For intermediate lengths ($L = 60\mu\text{m}$ to $80\mu\text{m}$) the stiffening is modest (a few to ~10% reduction), whereas for longer overhangs ($L \geq 100\mu\text{m}$) the effect becomes pronounced and can reduce the initial deflection by 15% to 40% at 20000 RPM.

To provide an absolute sense of scale, the reduction factor in (D.6) can be combined directly with the initial-deflection tables in Section D.2. For example, for $L = 150\mu\text{m}$ and $\Delta\sigma/d = 300\text{ MPa}\mu\text{m}^{-1}$, Table D.1 gives $h_0 \approx 16.1\mu\text{m}$; with $k_{\text{eff}}/k_0 = 1.60$ at 20000 RPM (Table D.4) this reduces to $h(\Omega) \approx h_0/1.60 \approx 10.1\mu\text{m}$. As a

Table D.4: Centrifugal stiffening factor k_{eff}/k_0 and corresponding relative tip-deflection reduction $\Delta h/h_0$ at $N = 20\,000$ RPM for the a-SiC design case ($w = 10\,\mu\text{m}$, $d = 96\,\text{nm}$, $r_o = 18\,\text{mm}$).

| L (μm) | k_{eff}/k_0 at 20 000 RPM | $\Delta h/h_0$ (%) |
|-----------------------|------------------------------------|--------------------|
| 5 | 1.00 | 0.0 |
| 10 | 1.00 | 0.0 |
| 40 | 1.01 | 1.1 |
| 60 | 1.04 | 3.7 |
| 80 | 1.09 | 8.3 |
| 100 | 1.18 | 15.1 |
| 120 | 1.31 | 23.4 |
| 150 | 1.60 | 37.4 |

D

second example, for $L = 60\,\mu\text{m}$ and $\Delta\sigma/d = 200\,\text{MPa}\mu\text{m}^{-1}$, $h_0 \approx 1.71\,\mu\text{m}$ and the reduction at 20 000 RPM is small ($k_{\text{eff}}/k_0 = 1.04$), giving $h(\Omega) \approx 1.65\,\mu\text{m}$.

Fig. D.3 summarises the predicted speed dependence by plotting k_{eff}/k_0 versus RPM for the considered overhang lengths. The corresponding deflection ratios follow directly from (D.6). In particular, the expected Ω^2 scaling appears as a strong increase of k_{eff}/k_0 with speed for long overhangs, while short overhangs remain nearly unaffected even at high RPM.

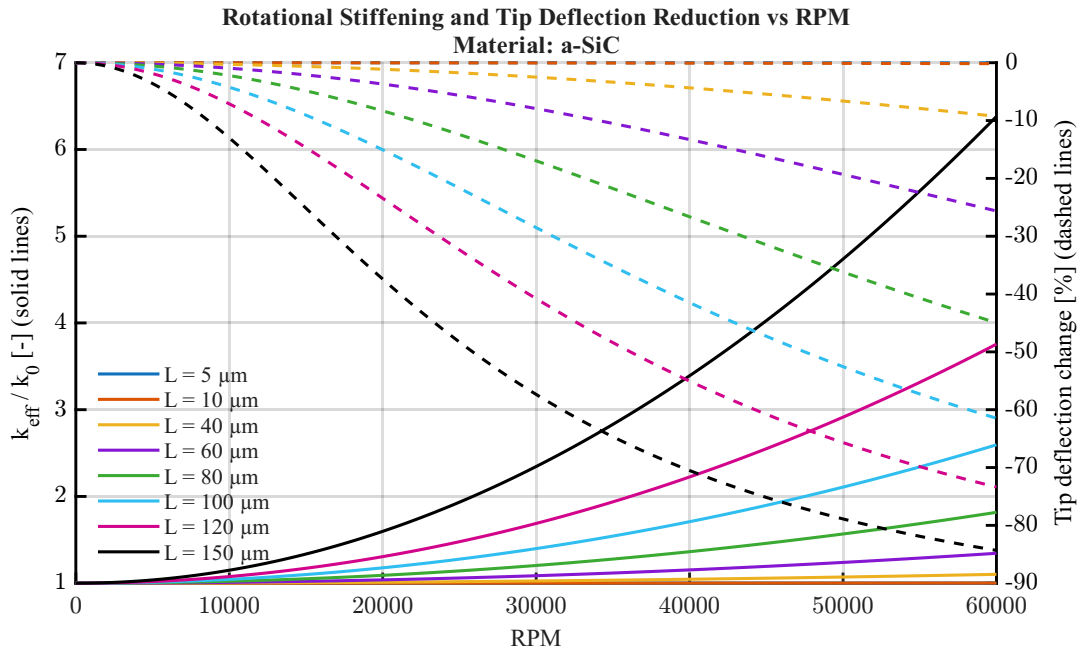


Figure D.3: Predicted centrifugal stiffening k_{eff}/k_0 versus rotational speed for the a-SiC design case ($w = 10\,\mu\text{m}$, $d = 96\,\text{nm}$, $r_o = 18\,\text{mm}$). The equivalent deflection ratio is $h(\Omega)/h_0 = 1/(k_{\text{eff}}/k_0)$, and the relative reduction is $\Delta h/h_0 = 1 - 1/(k_{\text{eff}}/k_0)$.

D.4.2. SiO₂ COMPARISON CASE (NUMERICAL ANCHOR)

For the thicker SiO₂ edge used in subsequent etching trials ($d = 4\,\mu\text{m}$), the same model predicts negligible centrifugal stiffening because $\alpha \propto A/I \propto 1/d^2$ in (D.7). As a numeric anchor, taking $E = 70\,\text{GPa}$, $\rho = 2300\,\text{kgm}^{-3}$, $w = 10\,\mu\text{m}$, $d = 4\,\mu\text{m}$, $r_o = 18\,\text{mm}$, $L = 150\,\mu\text{m}$, and $N = 60\,000$ RPM yields $k_{\text{eff}}/k_0 \approx 1.007$ (i.e. $\Delta h/h_0 \approx 0.7\%$). This confirms that rotational stiffening is an important mechanism for ultra-thin film edges, but becomes negligible for the much thicker SiO₂ case over the practically relevant undercut lengths studied in this thesis.

E

RAYLEIGH–RITZ DERIVATION

This appendix provides the derivation of the closed-form centrifugal-stiffening coefficient α used in Section 4.2.2 and Appendix D. Starting from a one-mode Rayleigh–Ritz (energy) formulation for a centrifugally tensioned Euler–Bernoulli cantilever, the result yields $k_{\text{eff}}/k_0 = 1 + \alpha$ and therefore $h(\Omega)/h_0 = 1/(1 + \alpha)$.

E.1. ASSUMPTIONS

Simplifying assumptions are made to obtain an approximate, closed-form solution from an otherwise non-linear boundary-value problem. The assumptions are:

- **Constant cross-section**

The ultra-thin cutting edge is modelled as a cantilever beam with a rectangular cross-section, for which the cross-sectional area $A(x) = A = wd$ and the area moment of inertia $I = \frac{wd^3}{12}$ are constant along the beam. Here, w is the in-plane width of the beam and d is the film thickness in the bending direction.

- **Euler–Bernoulli beam theory**

Cross-sections remain plane and normal to the neutral axis during deformation. The effects of shear deformation and rotational inertia are neglected. This simplification is appropriate for long and slender beams, for which $L \gg d$ [61].

- **Linear elasticity**

Following previous work, the a-SiC is modelled as a linearly elastic and isotropic solid in the small-strain regime [60], [62]. Young’s modulus E and density ρ are constant along the beam.

- **Constant angular velocity**

The beam rotates at a constant angular velocity Ω (in rad/s). The relation between Ω and the rotational speed N in revolutions per minute (RPM) is

$$\Omega = \frac{2\pi}{60} N. \quad (\text{E.1})$$

- **Representation of initial shape**

The initial deflection is represented by a single trial function (Section E.3) corresponding to the classical cubic deflection curve of a prismatic cantilever under a tip load. This shape enforces the essential clamped-end boundary conditions exactly and satisfies the zero-moment condition at the free end. The initial tip deflection is denoted by $h_0 = w(L)$, and the associated curvature scale is consistent with the small-deflection curvature estimate $R \approx L^2/(2h_0)$.

- **Neglect of Coriolis forces**

The configuration considered here is quasi-static. Only centrifugal forces are incorporated, Coriolis and gyroscopic effects are neglected [63].

- **Uncoupled motion**

Axial deformation is not solved for as an independent degree of freedom. Instead, the axial force distri-

bution $F_c(x)$ induced by centrifugal loading is computed on the undeformed geometry and treated as a known distribution. Its influence on bending is incorporated through a geometric-stiffness term in the potential energy (Section E.4), consistent with standard beam-column and rotating-beam theory [64], [65].

E.2. CENTRIFUGAL FORCE DISTRIBUTION

Under rotation at angular velocity Ω , each infinitesimal mass element along the beam experiences a centrifugal acceleration, leading to an axial force distribution $F_c(x)$. Following [63], the time-independent axial force at position x is

$$F_c(x) = \int_x^L \rho A \Omega^2 (R_E + \xi) d\xi, \quad (\text{E.2})$$

where R_E is the distance from the rotation axis to the clamped end and ξ is the integration variable. Evaluating (E.2) gives

$$F_c(x) = \rho A \Omega^2 \left[R_E(L-x) + \frac{1}{2}(L^2 - x^2) \right]. \quad (\text{E.3})$$

The maximal axial tension occurs at the root ($x=0$). This force is computed on the undeformed geometry. Small changes in radius due to transverse deflection are neglected.

E.3. REPRESENTING THE INITIAL CURVE

The initial deflected shape is approximated by the classical static solution of a prismatic cantilever beam of length L with a concentrated load P applied at the free end [55]:

$$w(x) = \frac{Px^2(3L-x)}{6EI}, \quad w_{\max} = w(L) = \frac{PL^3}{3EI}. \quad (\text{E.4})$$

A dimensionless shape function with unit tip deflection is then defined as

$$\phi(x) = \frac{w(x)}{w(L)} = \frac{x^2(3L-x)}{2L^3}. \quad (\text{E.5})$$

Its derivatives are

$$\phi'(x) = \frac{3(2Lx - x^2)}{2L^3}, \quad (\text{E.6})$$

$$\phi''(x) = \frac{3(L-x)}{L^3}, \quad (\text{E.7})$$

$$\phi'''(x) = -\frac{3}{L^3}. \quad (\text{E.8})$$

The shape function satisfies the essential boundary conditions at the clamped end and the zero-moment condition at the free end:

- $\phi(0) = 0$ (no deflection at the clamped end),
- $\phi'(0) = 0$ (zero slope at the clamped end),
- $\phi''(L) = 0$ (zero bending moment at the free end).

The zero-shear condition at the free end is not enforced ($\phi'''(L) \neq 0$). For Rayleigh–Ritz, the essential (displacement and slope) boundary conditions must be satisfied exactly; moment and shear at the free end are natural boundary conditions that need not be enforced by the shape function, although doing so can improve accuracy.

The transverse deflection is represented as

$$w(x) = h(\Omega) \phi(x), \quad (\text{E.9})$$

where $h(\Omega)$ is the scalar tip-deflection amplitude. At zero rotation $h(\Omega=0) = h_0$. At finite rotational speed $h(\Omega)$ is reduced due to centrifugal stiffening.

E.4. RAYLEIGH–RITZ ENERGY BALANCE

Let Q be the generalised transverse load producing the tip deflection h (e.g. due to residual stress). Its corresponding external work is $W = Qh$.

With the assumed shape (E.9), the bending strain energy is [64]

$$\begin{aligned} U_{\text{bending}} &= \frac{1}{2} \int_0^L EI w''(x)^2 dx = \frac{1}{2} \int_0^L EI (h\phi''(x))^2 dx \\ &= \frac{1}{2} EI h^2 \underbrace{\int_0^L \phi''(x)^2 dx}_{I_1}. \end{aligned} \quad (\text{E.10})$$

The additional strain energy associated with bending a beam under axial tension $F_c(x)$ (geometric stiffness) is [64]

$$\begin{aligned} U_{\text{axial}} &= \frac{1}{2} \int_0^L F_c(x) w'(x)^2 dx = \frac{1}{2} \int_0^L F_c(x) h^2 \phi'(x)^2 dx \\ &= \frac{1}{2} h^2 \underbrace{\int_0^L F_c(x) \phi'(x)^2 dx}_{I_2}. \end{aligned} \quad (\text{E.11})$$

The total potential energy then becomes

$$\Pi(h) = U_{\text{bending}} + U_{\text{axial}} - Qh = \frac{1}{2} h^2 (EII_1 + I_2) - Qh. \quad (\text{E.12})$$

For static equilibrium, Π must be stationary with respect to h :

$$\frac{\partial \Pi}{\partial h} = 0.$$

This yields

$$h(EII_1 + I_2) - Q = 0 \quad \Rightarrow \quad h(\Omega) = \frac{Q}{EII_1 + I_2}.$$

For zero rotation ($\Omega = 0$), $F_c(x) = 0 \Rightarrow I_2 = 0$, hence

$$h_0 = \frac{Q}{EII_1} \quad \Rightarrow \quad Q = h_0 EII_1. \quad (\text{E.13})$$

Substituting Q into the general expression gives the rotating tip deflection

$$h(\Omega) = \frac{h_0 EII_1}{EII_1 + I_2} = h_0 \left(\frac{1}{1 + \frac{I_2}{EII_1}} \right). \quad (\text{E.14})$$

E.5. EVALUATION OF THE INTEGRALS I_1 AND I_2

E.5.1. INTEGRAL I_1

Using (E.7),

$$\phi''(x) = \frac{3(L-x)}{L^3}, \quad (\text{E.15})$$

so that

$$\begin{aligned} I_1 &= \int_0^L \phi''(x)^2 dx = \left(\frac{3}{L^3} \right)^2 \int_0^L (L-x)^2 dx \\ &= \frac{9}{L^6} \left[\frac{x^3}{3} - Lx^2 + L^2x \right]_0^L = \frac{9}{L^6} \left(\frac{L^3}{3} \right) = \frac{3}{L^3}. \end{aligned} \quad (\text{E.16})$$

E.5.2. INTEGRAL I_2

Substituting (E.3) and (E.6) into I_2 gives

$$\begin{aligned} I_2 &= \int_0^L F_c(x) \phi'(x)^2 dx \\ &= \int_0^L \rho A \Omega^2 \left[R_E(L-x) + \frac{1}{2}(L^2 - x^2) \right] \left(\frac{3(2Lx - x^2)}{2L^3} \right)^2 dx. \end{aligned} \quad (\text{E.17})$$

Carrying out the integration yields

$$I_2 = \frac{3}{280} \rho A \Omega^2 (27L + 35R_E). \quad (\text{E.18})$$

The closed-form expression in (E.18) was verified symbolically in MATLAB.

E.5.3. FINAL EXPRESSION

Substituting I_1 and I_2 into (E.14) gives

$$\frac{I_2}{EI I_1} = \frac{\frac{3}{280} \rho A \Omega^2 (27L + 35R_E)}{EI \cdot \frac{3}{L^3}} = \frac{\rho A \Omega^2 L^3 (27L + 35R_E)}{280EI}. \quad (\text{E.19})$$

The resulting estimate for the tip deflection at angular speed Ω is therefore

$$h(\Omega) = \frac{h_0}{1 + \frac{\rho A \Omega^2 L^3 (27L + 35R_E)}{280EI}}. \quad (\text{E.20})$$

EQUIVALENT STIFFNESS INTERPRETATION

From the Rayleigh–Ritz expression for the potential energy (Eq. (E.12)),

$$\Pi(h) = \frac{1}{2} h^2 (EI I_1 + I_2) - Qh,$$

the equivalent bending stiffness can be expressed as

$$k_{\text{eff}}(L, \Omega) = EI I_1(L) + I_2(L, \Omega). \quad (\text{E.21})$$

At zero rotational speed, $I_2 = 0$ and the initial bending stiffness becomes

$$k_0(L) = k_{\text{eff}}(L, 0) = EI I_1(L) = 3EI/L^3. \quad (\text{E.22})$$

Under rotation, the stiffness increases due to centrifugal tension,

$$k_{\text{eff}}(L, \Omega) = k_0(L) + \frac{3}{280} \rho A \Omega^2 (27L + 35R_E), \quad (\text{E.23})$$

so that the stiffening factor is

$$\frac{k_{\text{eff}}(L, \Omega)}{k_0(L)} = 1 + \frac{\rho A \Omega^2 L^3 (27L + 35R_E)}{280EI}. \quad (\text{E.24})$$

Using $Q = k_0 h_0 = k_{\text{eff}} h(\Omega)$ it follows directly that

$$\frac{h(\Omega)}{h_0} = \frac{k_0(L)}{k_{\text{eff}}(L, \Omega)},$$

which is consistent with Eq. (E.20).

F

MATLAB CODE

This appendix provides the MATLAB script used to numerically evaluate the closed-form expressions for (i) static tip deflection due to a through-thickness stress gradient, (ii) the reduction in deflection under centrifugal stiffening during rotation, and (iii) the effective stiffness increase with rotational speed, as reported in the bending and stiffening appendix. To reproduce different material cases, edit the USER INPUT block (in particular the material preset values for E , ρ , and film thickness d) and rerun the script. The script writes a single summary .xlsx file containing the input parameters and the computed tables.

```
%% Tip deflection of rotating ultra-thin cantilever under centrifugal
    stiffening
% Computes:
% - Initial tip deflection h0 from a stress gradient (Delta sigma / d)
% - Dimensionless initial deflection (h0/L)*100%
% - Tip deflection during rotation h(Omega)
% - Percentage change in tip deflection due to rotation
% - Stiffness increase: k0(L), keff(L,Omega), and keff/k0
%
% Formulae correspond to the bending and rotational stiffening appendix
    in the thesis.

clear; clc;

%% ----- USER INPUT -----
% Output directory for summary .xlsx file (relative by default)
outputDir = fullfile(pwd, "Summaries"); % Change if desired

% --- Material preset (edit as needed) ---
% a-SiC preset: E_GPa = 210, rho = 2900, d_um = 0.096
% SiO2 preset: E_GPa = 70, rho = 2300, d_um = 4
material = 'a-SiC'; % Material name for run ID
E_GPa = 210; % [GPa] Young's modulus
rho = 2900; % [kg/m^3] density

% Geometry (film / beam)
w_um = 10; % [um] beam width
d_um = 0.096; % [um] beam thickness (film thickness)
Ro_m = 0.018; % [m] outer radius (distance from rotation
    axis to free end location)

% Rotation
RPM = 20000; % [rpm] rotational speed
```

```

% Stress gradient levels and beam lengths to evaluate
stressGrad_MPa_per_um = [50 100 150 200 250 300];           % [MPa/um]
L_um_vec              = [5 10 40 60 80 100 120 150];         % [um]

%% ----- UNIT CONVERSIONS TO SI (m, Pa, s, kg) -----
E    = E_GPa * 1e9;           % [Pa]
w    = w_um * 1e-6;          % [m]
d    = d_um * 1e-6;          % [m]
L_m  = L_um_vec(:) * 1e-6;    % [m] column vector of beam lengths

% Cross-sectional area and second moment of area (bending about
  thickness d)
A    = w * d;                 % [m^2]
I    = w * d^3 / 12;          % [m^4]

% Stress gradient (Delta sigma / d): MPa/um -> Pa/m
% 1 MPa = 1e6 Pa, 1 um = 1e-6 m, so factor = 1e12
stressGrad_SI = stressGrad_MPa_per_um * 1e12;   % [Pa/m], row vector

% Angular velocity [rad/s]
Omega = 2*pi * RPM / 60;      % [rad/s]

%% ----- INITIAL TIP DEFLECTION h0 FROM STRESS GRADIENT -----
% h0 = (Delta sigma / d) * L^2 / (2E)
[L_grid, grad_grid] = ndgrid(L_m, stressGrad_SI); % [m], [Pa/m]
h0_m = grad_grid .* L_grid.^2 ./ (2 * E);         % [m]
h0_nm = h0_m * 1e9;                               % [nm]

%% ----- TIP DEFLECTION DURING ROTATION h(Omega) -----
% h(Omega) = h0 / (1 + alpha)
% alpha = (rho * A * Omega^2 * L^3 * (27L + 35 R_E)) / (280 * E * I)
% with R_E = Ro - L

RE = Ro_m - L_m; % [m]
alpha_L = rho * A * Omega^2 .* L_m.^3 .* (27*L_m + 35*RE) ...
          / (280 * E * I); % [-], column vector

% Broadcast alpha_L across all stressGradient columns
h_nm = h0_nm ./ (1 + alpha_L); % [nm], same size as
  h0_nm

%% ----- PERCENTAGE CHANGE OF TIP DEFLECTION -----
% delta[%] = 100 * (h(Omega) - h0) / h0
delta_pct = 100 * (h_nm - h0_nm) ./ h0_nm;

%% ----- STIFFNESS (k0 and keff) AS A FUNCTION OF L AND RPM
  -----
% k0(L) = 3 E I / L^3
% keff = k0 + I2
% I2 = (3/280) * rho * A * Omega^2 * (27 L + 35 R_E)

k0_N_per_m = 3 * E * I ./ (L_m.^3); % [N/m]
I2_N_per_m = (3/280) * rho * A * Omega^2 .* (27*L_m + 35*RE); % [N/m]
]

```

```

keff_N_per_m = k0_N_per_m + I2_N_per_m; % [N/m
]
stiffening_factor = keff_N_per_m ./ k0_N_per_m; % [-]

StiffTable = table(L_um_vec(:), k0_N_per_m, keff_N_per_m,
    stiffening_factor, ...
    'VariableNames', {'L_um', 'k0_N_per_m', 'keff_N_per_m'
    , 'keff_over_k0'});

%% ----- BUILD RESULT TABLE -----
nGrad = numel(stressGrad_MPa_per_um);

% First column: beam lengths in um
data = L_um_vec(:);

varNames = cell(1, 1 + 3*nGrad);
varNames{1} = 'L_um';

for j = 1:nGrad
    g = stressGrad_MPa_per_um(j);

    data = [data, h0_nm(:, j), h_nm(:, j), delta_pct(:, j)];

    varNames{3*j - 1} = sprintf('h0_%d_nm', g);
    varNames{3*j} = sprintf('h_%d_nm', g);
    varNames{3*j + 1} = sprintf('delta_%d_pct', g);
end

T = array2table(data, 'VariableNames', varNames);

%% ----- BUILD INITIAL DEFLECTION-ONLY TABLE (h0 AND h0/L*100%)
-----
h0_over_L_pct = (h0_m ./ L_m) * 100; % [%], size [nL x nGrad]

data_h0 = L_um_vec(:);
varNames_h0 = cell(1, 1 + 2*nGrad);
varNames_h0{1} = 'L_um';

for j = 1:nGrad
    g = stressGrad_MPa_per_um(j);

    data_h0 = [data_h0, h0_nm(:, j), h0_over_L_pct(:, j)];

    varNames_h0{2*j} = sprintf('h0_%d_nm', g);
    varNames_h0{2*j+1} = sprintf('h0_over_L_%d_pct', g);
end

h0Table = array2table(data_h0, 'VariableNames', varNames_h0);

%% ----- DISPLAY INPUT PARAMETERS AND TABLE -----
fprintf('Input parameters used:\n');
fprintf(' RPM = %.0f rpm\n', RPM);
fprintf(' E = %.3g GPa\n', E_GPa);
fprintf(' rho = %.3f kg/m^3\n', rho);
fprintf(' w = %.3f um\n', w_um);
fprintf(' d = %.3f um\n', d_um);

```

```

fprintf(' Ro           = %.6f m\n', Ro_m);
fprintf('\nStress gradient levels [MPa/um]:\n ');
fprintf('%.0f ', stressGrad_MPa_per_um); fprintf('\n');
fprintf('\nBeam lengths L [um]:\n ');
fprintf('%.0f ', L_um_vec); fprintf('\n\n');

disp('Table of tip deflections (h0: static, h: rotating) and percentage
change:');
disp(T);

%% ----- BUILD RUN ID AND SAVE SUMMARY TO XLSX -----
runID_raw = sprintf('%s_w%.3fum_d%.3fum_%drpm', ...
                    material, w_um, d_um, RPM);
runID_safe = regexp(lower(runID_raw), '[^\w]', '_');
xlsxFileName = sprintf('tipDeflectionSummary_%s.xlsx', runID_safe);

if ~exist(outputDir, 'dir')
    mkdir(outputDir);
end
xlsxFilePath = fullfile(outputDir, xlsxFileName);

%% Build an "Input" sheet with parameters
paramNames = { ...
    'RunID'; 'Date'; 'Material'; 'RPM'; 'E'; 'rho'; 'w'; 'd'; 'Ro'; ...
    'StressGradientLevels'; 'L_values'};

paramValues = { ...
    runID_raw; datestr(now, 0); material; sprintf('%.0f rpm', RPM); ...
    sprintf('%.3g GPa', E_GPa); sprintf('%.3f kg/m^3', rho); ...
    sprintf('%.3f um', w_um); sprintf('%.3f um', d_um); sprintf('%.6f m
', Ro_m); ...
    sprintf('%s MPa/um', num2str(stressGrad_MPa_per_um)); ...
    sprintf('%s um', num2str(L_um_vec))};

paramUnits = { ...
    ''; ''; ''; 'rpm'; 'GPa'; 'kg/m^3'; 'um'; 'um'; 'm'; 'MPa/um'; 'um'
};

InputTable = table(paramNames, paramValues, paramUnits, ...
    'VariableNames', {'Parameter', 'Value', 'Unit'});

writetable(InputTable, xlsxFilePath, 'Sheet', 'Input', '
WriteVariableNames', true);
writetable(h0Table, xlsxFilePath, 'Sheet', 'InitialDeflection', '
WriteVariableNames', true);
writetable(T, xlsxFilePath, 'Sheet', 'Results', '
WriteVariableNames', true);
writetable(StiffTable, xlsxFilePath, 'Sheet', 'Stiffness', '
WriteVariableNames', true);

fprintf('Summary written to Excel file:\n %s\n', xlsxFilePath);

```

G

SUPPLEMENTARY FIGURES FOR CHALLENGE

4

G.1. IMAGING TRIALS FOR CLEARANCE INDICATION AT THE ROTATING RIM

Imaging trials were performed to assess whether the rotating rim could be visually “frozen” (i.e. rendered with negligible motion blur) to improve observation of the approach and contact region. At the target rim speed $v \approx 30$ m/s (at $r = 18$ mm), continuous illumination cannot freeze the rim at practical exposure settings. Even at the maximum shutter speed (1/8000 s), the apparent motion during exposure is $\Delta x = vt \approx 3.8$ mm. Consequently, the camera view was used only as a *qualitative clearance indication* during stepwise approach, rather than to resolve edge features in-flight.

Strobed illumination was also evaluated. In principle, sufficiently short light pulses can reduce motion blur. However, in this work the strobe could not be phase-locked to the rotor and the spindle speed exhibited small drift. As a result, the apparent edge position drifted over time. In addition, during rolling-shutter video acquisition, pulsed illumination introduced banding/artefacts that reduced the reliability of the view. For these reasons, strobed imaging was not used for the reported cut demonstrations. Fig. G.1 summarises representative trials.

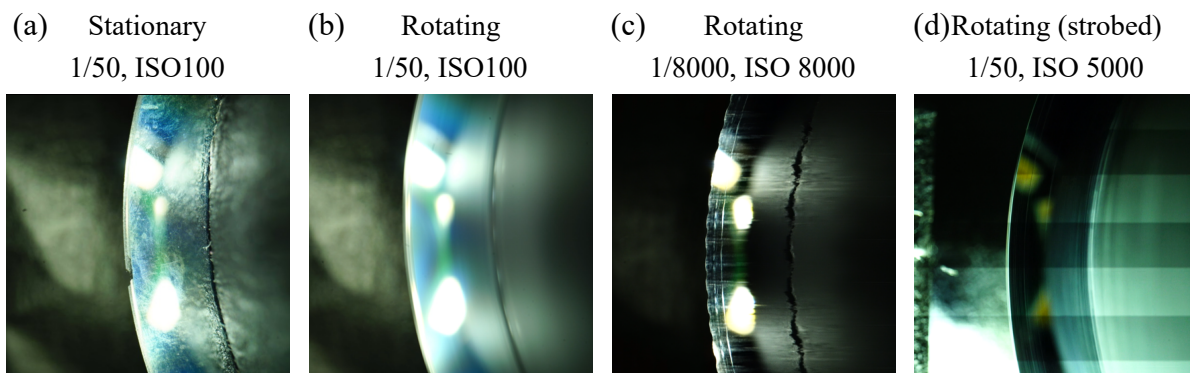


Figure G.1: Representative imaging trials near the cutting edge used to assess feasibility of reducing motion blur for approach monitoring: (a) stationary rim under continuous illumination, (b) rotating rim under continuous illumination (motion blur but stable view for qualitative clearance), (c) fast-shutter trial at high ISO (rolling-shutter distortion and residual blur), and (d) strobed illumination without phase-lock (apparent drift due to phase slip and banding/artefacts under rolling-shutter video). These observations motivated the use of imaging only for qualitative clearance indication during stepwise approach.

G.2. SAMPLE HOLDER FOR COMPLIANT ENGAGEMENT (THREAD TESTS)

To preserve the axial camera view during approach, the compliant sample (thread) must be held without obscuring the line-of-sight to the cutting edge. A 3D-printed sample holder was therefore designed to clamp

the sample in a diagonal orientation between two arms (Fig. G.2). This geometry keeps the region of interest visible when the sample is brought close to the blade (cf. Fig. 5.1(b)).

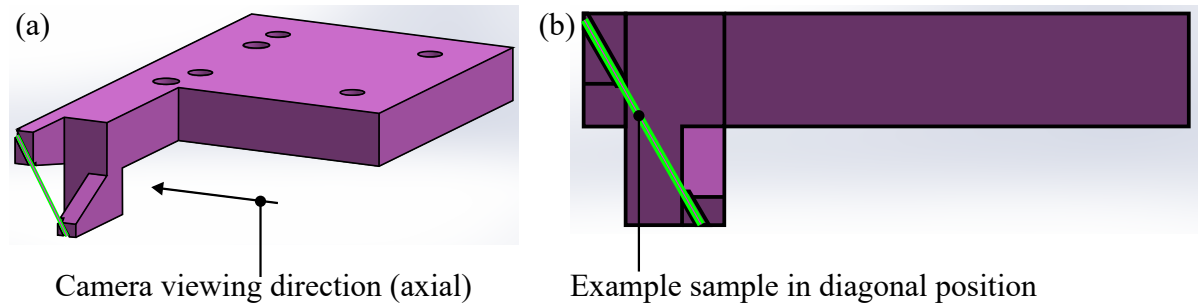


Figure G.2: 3D-printed sample holder used for the thread cut demonstrations. (a) Isometric CAD view of the holder geometry and mounting features (camera direction indicated). (b) Front view illustrating the diagonal thread sample orientation (green) used to maintain an unobstructed axial line-of-sight to the blade–sample contact region during approach (cf. Fig. 5.1).

G.3. SEWING-THREAD RESPONSE USED AS A QUALITATIVE CONTACT INDICATOR

As discussed in Section G.1, the rotating rim is strongly motion blurred at the target rim speed and first contact cannot be localised reliably from the camera view alone. For the sewing-thread tests (Tests 2–4), the thread was therefore used as a practical qualitative indicator of engagement: when the released rim contacts and severs a subset of fibres under tension, visible fibre separation occurs and can be observed in the camera view.

Fig. G.3 shows representative thread conditions prior to cutting and after partial cutting steps. The images support the qualitative interpretation used in the main text: visible fibre separation provided an immediate indication that engagement had occurred. No quantitative cut metrics are inferred from these images.




Figure G.3: Optical micrographs of the polyester sewing thread used as a qualitative contact indicator during the initial cut demonstrations. (a) Higher-magnification view of intact fibres prior to cutting. (b) Thread after Test 2, showing local fibre separation consistent with partial severing. (c) Thread after Test 3, showing a larger separated region consistent with further severing. Scale bars: 20 µm (a) and 100 µm (b–c).

G.4. POST-COMPLIANT SAMPLE TEST MICROSCOPY AND INTERPRETATION

Fig. G.4 summarises Keyence VHX-6000 observations before (Fig. G.4(a,b)) and after the thread tests (Fig. G.4(c–f)). Two observations from the front-side rim are particularly relevant. First, Fig. G.4(b) and Fig. G.4(c) show the same pre-existing local defect region (extent $\approx 902\mu\text{m}$ along the rim) before and after the cutting tests, respectively. The geometry and extent appear unchanged, indicating that this specific front-side sector did not experience detectable contact. Second, other locations on the front-side rim show large breakaway features (Fig. G.4(d)). These features were not present in pre-test inspection and were not observed during intermediate inspection between Test 3 and Test 4, indicating that the damage occurred during or after the final cut attempt.

Importantly, post-test microscopy itself can introduce damage to fragile released edges. During inspection it was observed that incidental contact between the rim and the microscope stage during placement (i.e.

imperfectly flat seating causing a local edge–stage touch-down) could introduce additional edge damage to the front-side rim. The rear-side (thicker) rim appeared less susceptible to this placement-induced damage. For future inspection, the blade should be supported on a central pedestal with radius smaller than the blade radius, such that the rims cannot contact a flat support surface during placement.

In contrast to the front-side observations, the rear-side (thicker) rim exhibits local regions of slight wear (Fig. G.4(e)) while other inspected regions appear unchanged (Fig. G.4(f)) and no chipping of comparable severity was observed on the rear-side rim. The spatially localised nature of this wear is consistent with intermittent contact occurring only over a limited circumferential sector. One plausible explanation is that operational displacement (radial and/or axial) causes the rim to intermittently enter contact such that specific circumferential regions accumulate the majority of wear. Overall, the microscopy supports the qualitative conclusion that the released rims can survive rotation at the required rim speed and can perform a first cut demonstration on a soft, tensioned sample. However, the thread tests do not allow unambiguous attribution of cutting action to the front-side $\sim 4\mu\text{m}$ rim versus the rear-side $\sim 10\mu\text{m}$ rim.

G.5. POST-TEST 5 MICROSCOPY AND INTERPRETATION

Fig. G.5(a,b) shows that the main damaged band on the polymer has a measured width of $547\mu\text{m}$. This value is comparable to the total blade stack thickness ($\sim 4\mu\text{m} + 525\mu\text{m} + 10\mu\text{m} \approx 539\mu\text{m}$), which is consistent with over-travel leading to contact of the Si sidewall and subsequent ploughing by the substrate, rather than cutting dominated by the ultra-thin released edge(s).

At the far end of the track (Fig. G.5(c)), the initial feature is locally $\sim 10\mu\text{m}$ wide. This is compatible with initial engagement by the rear-side released oxide thickness prior to sidewall contact. However, it is not definitive, since local chipping, debris drag, or partial rim breakaway can produce similar apparent widths. Overall, this trial underscores the need for a reliable first-engagement indicator and improved optical access to prevent over-travel in rigid-sample tests.

G.6. REPRESENTATIVE WLI MAP AND LINE PROFILE (TEST 6 KERF DEPTH)

White-light interferometry (WLI) was used to quantify the kerf depth produced in the controlled rigid-sample trial (Test 6). The depth metric reported in the main text corresponds to the *maximum trench depth relative to the local mean surface plane* outside the kerf (i.e. peak-to-valley depth after levelling to the surrounding surface). Fig. G.6 shows a representative height map and corresponding line profile across the kerf near its mid-length region. The map confirms a pronounced trench-like depression with a maximum depth of order $10\mu\text{m}$ to $15\mu\text{m}$, consistent with the repeated measurements reported in Section 5.3.

G.7. POST-TEST 6 MICROSCOPY AND INTERPRETATION

To support rim attribution for the controlled rigid-sample trial (Test 6), the blade was inspected after the test using optical microscopy. Fig. G.7 provides representative examples of the rear-side rim condition after the test. Two characteristic appearances were observed: (i) regions exhibiting slight local degradation, expressed as increased edge roughness and a less uniform rim outline, and (ii) regions showing no apparent change compared with intact pre-test appearance within the optical resolution.

The spatially localised nature of the degradation is consistent with intermittent contact occurring only over a limited circumferential sector, rather than circumferentially uniform wear. In combination with the single preserved kerf track on the ABS sample and the absence of substrate sidewall ploughing signatures in Test 6, these observations support the interpretation in the main text that the rear-side ($\sim 10\mu\text{m}$ thick) rim was the dominant cutter under the present conditions. The examples also illustrate that, despite local wear, large-scale circumferential rim loss did not occur during Test 6.





Figure G.4: Optical micrographs (Keyence VHX-6000) of the released SiO₂ rims before and after the tensioned sewing-thread tests (Tests 2–4), used to assess edge condition and potential contact localisation. (a–d) Front-side rim on the polished wafer side (thin oxide, ~ 4 μm): (a) representative intact region with nominal protrusion length ($L \approx 120 \mu\text{m}$); (b) pre-existing local defect region prior to cutting (extent $\approx 902 \mu\text{m}$ along the rim); (c) the same region after cutting, showing an unchanged profile within the inspection resolution; (d) example of a severe local breakaway observed after the final cut attempt (origin uncertain). (e–f) Rear-side rim on the unpolished wafer side (thicker oxide, ~ 10 μm): (e) example region showing slight local wear; (f) example region showing no apparent change. Scale bars: 100 μm (a–d) and 20 μm (e–f).



G

Figure G.5: Post-test optical micrographs of a rigid polymer trial (white LEGO/ABS, Test 5) at $N \approx 16.9$ krPM. Boxed regions indicate the fields of view shown at higher magnification (a→b, b→c). (a) Overview of the modified surface region. (b) Apparent damage-band width $w = 547 \mu\text{m}$, comparable to the blade stack thickness ($\sim 539 \mu\text{m}$) and consistent with over-travel with substrate sidewall involvement (potential polymer pile-up and projection bias). The separate pre-test mark resulted from a prior touch with a clean Si blade. (c) Detail near the track end showing a local initial width of $\sim 10 \mu\text{m}$, compatible with the rear-side rim thickness but not definitive. Scale bars: 500 μm (a,b) and 20 μm (c).

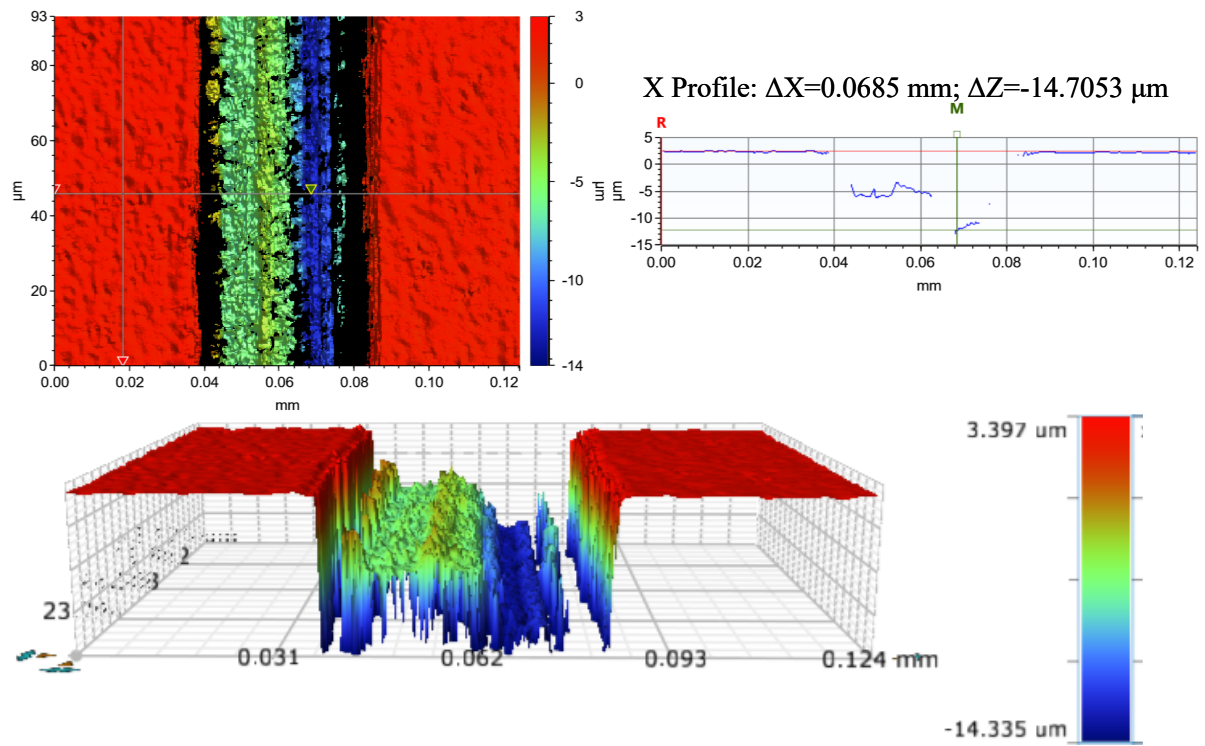


Figure G.6: Representative white-light interferometry (WLI) measurement of the Test 6 kerf on ABS. The top-left panel shows the height map of a region spanning the kerf; the top-right panel shows an example line profile across the trench. Depth is reported relative to the local mean surface plane outside the kerf. In this representative scan, the profile indicates a maximum trench depth of $\Delta z \approx 14.7 \mu\text{m}$ over a lateral span of $\Delta x \approx 68.5 \mu\text{m}$. The measurement supports the depth range reported in the main text for repeated scans near the kerf centre.




Figure G.7: Representative post-test microscopy of the rear-side rim after the controlled rigid-sample trial (Test 6). (a) Example region showing no apparent change within the inspection resolution. (b) Example region showing slight local degradation (increased edge roughness) over a limited sector. The coexistence of degraded and intact regions indicates spatially localised wear, consistent with intermittent contact limited to a circumferential sector. Scale bars: $50 \mu\text{m}$.

BIBLIOGRAPHY

- [1] S. Harmand et al., “3.3-million-year-old stone tools from lomekwi 3, west turkana, kenya,” *Nature*, vol. 521, no. 7552, pp. 310–315, May 20, 2015, ISSN: 0028-0836, 1476-4687. DOI: [10.1038/nature14464](https://doi.org/10.1038/nature14464). Accessed: Nov. 12, 2024. [Online]. Available: <https://www.nature.com/articles/nature14464>.
- [2] L. S. B. Leakey, P. V. Tobias, and J. R. Napier, “A new species of the genus homo from olduvai gorge,” *Nature*, vol. 202, no. 4927, pp. 7–9, Apr. 4, 1964, ISSN: 0028-0836, 1476-4687. DOI: [10.1038/202007a0](https://doi.org/10.1038/202007a0). Accessed: Nov. 12, 2024. [Online]. Available: <https://www.nature.com/articles/202007a0>.
- [3] A. Lucas and J. Harris, “Wood (foreign timber: Egyptian timber): Woodworking: Bark: Silicified wood: Charcoal,” in *Ancient Egyptian Materials and Industries*, Dover Publications, Apr. 30, 2012, p. 449, ISBN: 978-1-306-47501-3.
- [4] M. P. Bale, “Circular saw benches,” in *Woodworking Machinery: Its Rise, Progres and Construction*, London: Crosby Lockwood and Co., 1880, p. 6.
- [5] R. S. Woodbury, “The clock and instrument makers,” in *History of the Gear-Cutting Machine: A Historical Study in Geometry and Machines*, Cambridge, Massachusetts: The M.I.T. Press, 1964, pp. 45–50, ISBN: 978-0-262-73001-3.
- [6] N. Ball, “Circular saws and the history of technology,” *Bulletin of the Association for Preservation Technology*, vol. 7, no. 3, pp. 79–89, 1975, ISSN: 00449466. DOI: [10.2307/1493506](https://doi.org/10.2307/1493506). Accessed: Nov. 13, 2024. [Online]. Available: <https://www.jstor.org/stable/1493506?origin=crossref>.
- [7] V. Kocovic, D. Vukelic, S. Kostic, I. Bijelic, M. Prica, and B. Tadić, “Micro-cutting of holes by centrifugal force,” *The International Journal of Advanced Manufacturing Technology*, vol. 124, no. 5, pp. 1437–1455, Jan. 1, 2023, ISSN: 1433-3015. DOI: [10.1007/s00170-022-10581-w](https://doi.org/10.1007/s00170-022-10581-w). Accessed: Jan. 5, 2025. [Online]. Available: <https://doi.org/10.1007/s00170-022-10581-w>.
- [8] M. A. Lieberman and A. J. Lichtenberg, *Principles of Plasma Discharges and Materials Processing*, 1st ed. Wiley, Apr. 8, 2005, ISBN: 978-0-471-72001-0 978-0-471-72425-4. DOI: [10.1002/0471724254](https://doi.org/10.1002/0471724254). Accessed: Mar. 3, 2026. [Online]. Available: <https://onlinelibrary.wiley.com/doi/book/10.1002/0471724254>.
- [9] L. A. Giannuzzi and F. A. Stevie, Eds., *Introduction to Focused Ion Beams: Instrumentation, Theory, Techniques and Practice*, Boston, MA: Springer US, 2005, ISBN: 978-0-387-23116-7 978-0-387-23313-0. DOI: [10.1007/b101190](https://doi.org/10.1007/b101190). Accessed: Mar. 3, 2026. [Online]. Available: <http://link.springer.com/10.1007/b101190>.
- [10] E. Jüngling, S. Wilczek, T. Mussenbrock, M. Böke, and A. Von Keudell, “Plasma sheath tailoring by a magnetic field for three-dimensional plasma etching,” *Applied Physics Letters*, vol. 124, no. 7, p. 074 101, Feb. 12, 2024, ISSN: 0003-6951, 1077-3118. DOI: [10.1063/5.0187685](https://doi.org/10.1063/5.0187685). Accessed: Dec. 30, 2025. [Online]. Available: <https://pubs.aip.org/apl/article/124/7/074101/3262928/Plasma-sheath-tailoring-by-a-magnetic-field-for>.
- [11] Y. Zhao, X. Zhou, J. Zhang, S. Song, and Y. Zhao, “Improving plasma uniformity in the inductively coupled plasma by external magnetic field,” *Physics of Plasmas*, vol. 31, no. 8, p. 083 507, Aug. 1, 2024, ISSN: 1070-664X, 1089-7674. DOI: [10.1063/5.0214357](https://doi.org/10.1063/5.0214357). Accessed: Dec. 30, 2025. [Online]. Available: <https://pubs.aip.org/pop/article/31/8/083507/3305914/Improving-plasma-uniformity-in-the-inductively>.
- [12] O. Koplak, F. Maspero, F. Marson, M. Cocconcelli, A. Plaza, and R. Bertacco, “Effect of patterning on SmCo micromagnets suitable for integration in microsystems,” *Journal of Magnetism and Magnetic Materials*, vol. 605, p. 172 323, Sep. 1, 2024, ISSN: 0304-8853. DOI: [10.1016/j.jmmm.2024.172323](https://doi.org/10.1016/j.jmmm.2024.172323). Accessed: Mar. 10, 2026. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0304885324006140>.

- [13] T. Lisec, O. Behrmann, and B. Gojdka, "PowderMEMS—a generic microfabrication technology for integrated three-dimensional functional microstructures," *Micromachines*, vol. 13, no. 3, Feb. 28, 2022, ISSN: 2072-666X. DOI: [10.3390/mi13030398](https://doi.org/10.3390/mi13030398). Accessed: Mar. 10, 2026. [Online]. Available: <https://www.mdpi.com/2072-666X/13/3/398>.
- [14] Z. W. Xu et al., "Fabrication of micro DOE using micro tools shaped with focused ion beam," *Optics Express*, vol. 18, no. 8, pp. 8025–8032, Apr. 12, 2010, ISSN: 1094-4087. DOI: [10.1364/OE.18.008025](https://doi.org/10.1364/OE.18.008025). Accessed: Dec. 10, 2024. [Online]. Available: <https://opg.optica.org/oe/abstract.cfm?uri=oe-18-8-8025>.
- [15] H. T. Wang and W. B. Lee, "A study of cutting factors affecting the generation of functional hierarchical rib array structure in ultra-precision raster milling," *The International Journal of Advanced Manufacturing Technology*, vol. 86, no. 1, pp. 989–998, Sep. 2016, ISSN: 1433-3015. DOI: [10.1007/s00170-015-8210-z](https://doi.org/10.1007/s00170-015-8210-z). Accessed: Jan. 4, 2025. [Online]. Available: <https://doi.org/10.1007/s00170-015-8210-z>.
- [16] R. Huang, X. Zhang, W. K. Neo, A. S. Kumar, and K. Liu, "Ultra-precision machining of grayscale pixelated micro images on metal surface," *Precision Engineering*, vol. 52, pp. 211–220, Apr. 1, 2018, ISSN: 0141-6359. DOI: [10.1016/j.precisioneng.2017.12.009](https://doi.org/10.1016/j.precisioneng.2017.12.009). Accessed: Dec. 10, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0141635917305925>.
- [17] X. Zhang, K. Liu, V. Sunappan, and X. Shan, "Diamond micro engraving of gravure roller mould for roll-to-roll printing of fine line electronics," *Journal of Materials Processing Technology*, vol. 225, pp. 337–346, Nov. 1, 2015, ISSN: 0924-0136. DOI: [10.1016/j.jmatprotec.2015.05.032](https://doi.org/10.1016/j.jmatprotec.2015.05.032). Accessed: Dec. 10, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0924013615002381>.
- [18] R. Takigawa, E. Higurashi, T. Kawanishi, and T. Asano, "Lithium niobate ridged waveguides with smooth vertical sidewalls fabricated by an ultra-precision cutting method," *Optics Express*, vol. 22, no. 22, pp. 27 733–27 738, Nov. 3, 2014, ISSN: 1094-4087. DOI: [10.1364/OE.22.027733](https://doi.org/10.1364/OE.22.027733). Accessed: Dec. 10, 2024. [Online]. Available: <https://opg.optica.org/oe/abstract.cfm?uri=oe-22-22-27733>.
- [19] A.-K. Holthusen, O. Riemer, and E. Brinksmeier, "Material impact on diamond machining of diffractive optical structures for UV-application," *Journal of Manufacturing and Materials Processing*, vol. 2, no. 1, p. 15, Mar. 2018, ISSN: 2504-4494. DOI: [10.3390/jmmp2010015](https://doi.org/10.3390/jmmp2010015). Accessed: Jan. 3, 2025. [Online]. Available: <https://www.mdpi.com/2504-4494/2/1/15>.
- [20] T. Gietzelt and L. Eichhorn, "Herstellung von mikrostrukturen mit hohem aspektverhältnis durch mikrozerspannung," *Galvanotechnik*, vol. 89, no. 6, pp. 1504–1514, 2007.
- [21] J. Sun, X. Luo, W. Chang, J. M. Ritchie, J. Chien, and A. Lee, "Fabrication of periodic nanostructures by single-point diamond turning with focused ion beam built tool tips," *Journal of Micromechanics and Microengineering*, vol. 22, no. 11, p. 115014, Nov. 1, 2012, ISSN: 0960-1317, 1361-6439. DOI: [10.1088/0960-1317/22/11/115014](https://doi.org/10.1088/0960-1317/22/11/115014). Accessed: Dec. 10, 2024. [Online]. Available: <https://iopscience.iop.org/article/10.1088/0960-1317/22/11/115014>.
- [22] T. Schaller, L. Bohn, J. Mayer, and K. Schubert, "Microstructure grooves with a width of less than 50 μm cut with ground hard metal micro end mills," *Precision Engineering*, vol. 23, no. 4, pp. 229–235, Oct. 1, 1999, ISSN: 0141-6359. DOI: [10.1016/S0141-6359\(99\)00011-2](https://doi.org/10.1016/S0141-6359(99)00011-2). Accessed: Dec. 9, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0141635999000112>.
- [23] D. P. Adams, M. J. Vasile, G. L. Benavides, and A. N. Campbell, "Micromilling of metal alloys with focused ion beam-fabricated tools," *Journal of Precision Engineering*, Oct. 3, 2000. Accessed: Dec. 9, 2024. [Online]. Available: <https://www.osti.gov/biblio/764822>.
- [24] H. Hoffmeister and A. Wenda, "Investigations into the microgrinding technique for the production of microstructures in silicon and glass micro system technologies 98," in *6th International Conference on Micro Electro, Opto, Mechanical Systems and Components*, 1998, pp. 433–438.
- [25] P. Loiko et al., "Channel waveguide lasers in bulk tm:LiYF_4 produced by deep diamond-saw dicing," *Optics Express*, vol. 28, no. 18, pp. 26 676–26 689, Aug. 31, 2020, ISSN: 1094-4087. DOI: [10.1364/OE.400497](https://doi.org/10.1364/OE.400497). Accessed: Feb. 26, 2025. [Online]. Available: <https://opg.optica.org/oe/abstract.cfm?uri=oe-28-18-26676>.

- [26] E. Brinksmeier, Y. Mutlugünes, F. Klocke, J. C. Aurich, P. Shore, and H. Ohmori, "Ultra-precision grinding," *CIRP Annals*, vol. 59, no. 2, pp. 652–671, Jan. 1, 2010, ISSN: 0007-8506. DOI: [10.1016/j.cirp.2010.05.001](https://doi.org/10.1016/j.cirp.2010.05.001). Accessed: Dec. 4, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0007850610001885>.
- [27] J. Rönn et al., "Ultra-high on-chip optical gain in erbium-based hybrid slot waveguides," *Nature Communications*, vol. 10, no. 1, p. 432, Jan. 25, 2019, ISSN: 2041-1723. DOI: [10.1038/s41467-019-08369-w](https://doi.org/10.1038/s41467-019-08369-w). Accessed: Feb. 26, 2025. [Online]. Available: <https://www.nature.com/articles/s41467-019-08369-w>.
- [28] Z. Yuan, J. Hu, Q. Wen, K. Cheng, and P. Zheng, "Investigation on an innovative method for high-speed low-damage micro-cutting of CFRP composites with diamond dicing blades," *Materials*, vol. 11, p. 1974, Oct. 13, 2018. DOI: [10.3390/ma11101974](https://doi.org/10.3390/ma11101974).
- [29] Z. Yuan, A. Riaz, and B. s. Chohan, "Precision machining by dicing blades: A systematic review," *Machines*, vol. 11, no. 2, p. 259, Feb. 2023, ISSN: 2075-1702. DOI: [10.3390/machines11020259](https://doi.org/10.3390/machines11020259). Accessed: Sep. 27, 2024. [Online]. Available: <https://www.mdpi.com/2075-1702/11/2/259>.
- [30] X. Liu, R. E. DeVor, S. G. Kapoor, and K. F. Ehmann, "The mechanics of machining at the microscale: Assessment of the current state of the science," *Journal of Manufacturing Science and Engineering*, vol. 126, no. 4, pp. 666–678, Nov. 2004, ISSN: 1087-1357, 1528-8935. DOI: [10.1115/1.1813469](https://doi.org/10.1115/1.1813469). Accessed: Dec. 15, 2024. [Online]. Available: <https://asmedigitalcollection.asme.org/manufacturingscience/article/126/4/666/462128/The-Mechanics-of-Machining-at-the-Microscale>.
- [31] S. J. Zhang, S. To, S. J. Wang, and Z. W. Zhu, "A review of surface roughness generation in ultra-precision machining," *International Journal of Machine Tools and Manufacture*, vol. 91, pp. 76–95, Apr. 2015, ISSN: 0890-6955. DOI: [10.1016/j.ijmachtools.2015.02.001](https://doi.org/10.1016/j.ijmachtools.2015.02.001). Accessed: Dec. 15, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0890695515300092>.
- [32] S.-C. Kim, E.-S. Lee, N.-H. Kim, and H.-D. Jeong, "Machining characteristics on the ultra-precision dicing of silicon wafer," *The International Journal of Advanced Manufacturing Technology*, vol. 33, no. 7, pp. 662–667, Jul. 2007, ISSN: 0268-3768, 1433-3015. DOI: [10.1007/s00170-006-0499-1](https://doi.org/10.1007/s00170-006-0499-1). Accessed: Oct. 10, 2024. [Online]. Available: <https://link.springer.com/10.1007/s00170-006-0499-1>.
- [33] M. Li et al., "Ultrathin diamond blades for dicing single crystal SiC developed using a novel bonding method," *Journal of Manufacturing Processes*, vol. 84, pp. 88–99, Dec. 2022, ISSN: 1526-6125. DOI: [10.1016/j.jmapro.2022.09.058](https://doi.org/10.1016/j.jmapro.2022.09.058). Accessed: Dec. 15, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1526612522006788>.
- [34] T. He et al., "Investigation on the machining performance of copper-based diamond ultra-thin dicing blades manufactured by fused deposition modeling and sintering (FDMS)," *Tribology International*, vol. 187, p. 108702, Sep. 2023, ISSN: 0301-679X. DOI: [10.1016/j.triboint.2023.108702](https://doi.org/10.1016/j.triboint.2023.108702). Accessed: Dec. 15, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0301679X23004905>.
- [35] Y. Feng, K. Li, Z. Dou, Z. Zhang, and B. Guo, "High-speed dicing of SiC wafers with 0.048 mm diamond blades via rolling-slitting," *Materials*, vol. 15, no. 22, p. 8083, Nov. 15, 2022, ISSN: 1996-1944. DOI: [10.3390/ma15228083](https://doi.org/10.3390/ma15228083). Accessed: Dec. 15, 2024. [Online]. Available: <https://www.mdpi.com/1996-1944/15/22/8083>.
- [36] L. Zhang, H. Yu, and H. Ma, "Preparation and performance study of electroplated ni-w/diamond ultrathin dicing blades," *Materials Science in Semiconductor Processing*, vol. 175, p. 108280, Jun. 1, 2024, ISSN: 1369-8001. DOI: [10.1016/j.mssp.2024.108280](https://doi.org/10.1016/j.mssp.2024.108280). Accessed: Dec. 15, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1369800124001768>.
- [37] K. Matsumaru, A. Takata, and K. Ishizaki, "Advanced thin dicing blade for sapphire substrate," *Science and Technology of Advanced Materials*, 21st Century COE Program, Nagaoka University of Technology, Hybridized Materials with Super Functions 2005, vol. 6, no. 2, pp. 120–122, Mar. 1, 2005, ISSN: 1468-6996. DOI: [10.1016/j.stam.2004.11.002](https://doi.org/10.1016/j.stam.2004.11.002). Accessed: Dec. 4, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1468699605000033>.
- [38] H. Zhou, S. Qiu, Y. Huo, and N. Zhang, "High-speed dicing of silicon wafers conducted using ultrathin blades," *The International Journal of Advanced Manufacturing Technology*, vol. 66, no. 5, pp. 947–953, Jul. 19, 2012, ISSN: 1433-3015. DOI: [10.1007/s00170-012-4379-6](https://doi.org/10.1007/s00170-012-4379-6). Accessed: Dec. 15, 2024. [Online]. Available: <https://doi.org/10.1007/s00170-012-4379-6>.

- [39] S. P. Timošenko and J. N. Goodier, *Theory of elasticity* (Engineering societies monographs), 2. ed. New York: McGraw-Hill, 1951, 506 pp.
- [40] D. Roundy and M. L. Cohen, “Ideal strength of diamond, si, and ge,” *Physical Review B*, vol. 64, no. 21, p. 212 103, Nov. 12, 2001, ISSN: 0163-1829, 1095-3795. DOI: [10.1103/PhysRevB.64.212103](https://doi.org/10.1103/PhysRevB.64.212103). Accessed: Nov. 7, 2025. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.64.212103>.
- [41] D. Echizenya, H. Sakamoto, and K. Sasaki, “Effect of mechanical surface damage on silicon wafer strength,” *Procedia Engineering*, vol. 10, pp. 1440–1445, 2011, ISSN: 18777058. DOI: [10.1016/j.proeng.2011.04.239](https://doi.org/10.1016/j.proeng.2011.04.239). Accessed: Nov. 7, 2025. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S1877705811004279>.
- [42] J.-H. Woo et al., “Critical bending radius of thin single-crystalline silicon with dome and pyramid surface texturing,” *Scripta Materialia*, vol. 140, pp. 1–4, Nov. 2017, ISSN: 13596462. DOI: [10.1016/j.scriptamat.2017.06.047](https://doi.org/10.1016/j.scriptamat.2017.06.047). Accessed: Nov. 7, 2025. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S1359646217303597>.
- [43] Jie-Hua Zhao, J. Tellkamp, V. Gupta, and D. Edwards, “Experimental evaluations of the strength of silicon die by 3-point-bend versus ball-on-ring tests,” *IEEE Transactions on Electronics Packaging Manufacturing*, vol. 32, no. 4, pp. 248–255, Oct. 2009, ISSN: 1521-334X, 1558-0822. DOI: [10.1109/TEPM.2009.2028329](https://doi.org/10.1109/TEPM.2009.2028329). Accessed: Nov. 7, 2025. [Online]. Available: <http://ieeexplore.ieee.org/document/5256178/>.
- [44] J. Laconte, D. Flandre, and J.-P. Raskin, Eds., *Micromachined Thin-Film Sensors for SOI-CMOS Co-Integration*, SpringerLink Bücher, Boston, MA: Springer US, 2006, 292 pp., ISBN: 978-0-387-28842-0 978-0-387-28843-7. DOI: [10.1007/0-387-28843-0](https://doi.org/10.1007/0-387-28843-0).
- [45] HST Spindle Repair. “Accretech spindle repair and rebuild,” Accessed: Jan. 26, 2026. [Online]. Available: <https://hstspindles.com/portfolio/accretech-spindle-repair-and-rebuild/>.
- [46] Giorgio Technology Sales/Service. “Spindle rebuild, spindle repair,” Accessed: Jan. 26, 2026. [Online]. Available: <https://www.gtsaz.com/service/spindle-rebuild-repair.php>.
- [47] Air Bearings Limited. “DS060e dicing spindle,” Accessed: Jan. 26, 2026. [Online]. Available: <https://airbearings.co.uk/ds060e-dicing-spindle/>.
- [48] SKF Group, “Rolling bearings,” SKF Group, Catalogue PUB BU/P1 17000/1 EN, Oct. 2018. [Online]. Available: <https://www.skf.com/go/17000>.
- [49] LASEA, “Laser micromachining station,” Offer OFF190144 v1.1.
- [50] Y. Kilinc, U. Unal, and B. E. Alaca, “Residual stress gradients in electroplated nickel thin films,” *MICRO-ELECTRONIC ENGINEERING*, vol. 134, pp. 60–67, Jan. 1, 2015. DOI: [10.1016/j.mee.2015.01.042](https://doi.org/10.1016/j.mee.2015.01.042). Accessed: Nov. 25, 2025. [Online]. Available: <https://aperta.ulakbim.gov.tr/record/80981>.
- [51] P. Argyrakis, P. McNabb, A. J. Snell, and R. Cheung, “Relaxation of process induced surface stress in amorphous silicon carbide thin films using low energy ion bombardment,” *Applied Physics Letters*, vol. 89, no. 3, p. 034 101, Jul. 17, 2006, ISSN: 0003-6951. DOI: [10.1063/1.2222318](https://doi.org/10.1063/1.2222318). Accessed: Nov. 25, 2025. [Online]. Available: <https://doi.org/10.1063/1.2222318>.
- [52] W. Fang and J. A. Wickert, “Determining mean and gradient residual stresses in thin films using micromachined cantilevers,” *Journal of Micromechanics and Microengineering*, vol. 6, no. 3, p. 301, Sep. 1996, ISSN: 0960-1317. DOI: [10.1088/0960-1317/6/3/002](https://doi.org/10.1088/0960-1317/6/3/002). Accessed: Nov. 25, 2025. [Online]. Available: <https://doi.org/10.1088/0960-1317/6/3/002>.
- [53] C.-K. Chung, T.-Y. Lin, J.-G. Duh, and M.-Q. Tsai, “Global and local residual stress in silicon carbide films produced by plasma-enhanced chemical vapor deposition,” *Surface and Coatings Technology*, vol. 200, no. 16, pp. 4825–4834, Apr. 27, 2006, ISSN: 0257-8972. DOI: [10.1016/j.surfcoat.2005.04.048](https://doi.org/10.1016/j.surfcoat.2005.04.048). Accessed: Nov. 25, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0257897205005761>.
- [54] S. Maruthoor, A. Ajayakumar, T. Fuchs, O. Jakovlev, H. Reinecke, and J. Wilde, “Mechanical characterization of polycrystalline and amorphous silicon carbide thin films using bulge test,” *Journal of Microelectromechanical Systems*, vol. 22, no. 1, pp. 140–146, Feb. 2013, ISSN: 1941-0158. DOI: [10.1109/JMEMS.2012.2218577](https://doi.org/10.1109/JMEMS.2012.2218577). Accessed: Nov. 26, 2025. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/6328232>.

- [55] W. Beitz and K.-H. Küttner, Eds., *Dubbel Handbook of Mechanical Engineering*, London: Springer London, 1994, ISBN: 978-1-4471-3568-5. Accessed: Sep. 27, 2024. [Online]. Available: <https://link.springer.com/10.1007/978-1-4471-3566-1>.
- [56] S. Shamray, B. Azarhoushang, M. Paknejad, and A. Buechler, "Ductile-brittle transition mechanisms in micro-grinding of silicon nitride," *Ceramics International*, vol. 48, no. 23, pp. 34 987–34 998, Dec. 1, 2022, ISSN: 0272-8842. DOI: [10.1016/j.ceramint.2022.08.088](https://doi.org/10.1016/j.ceramint.2022.08.088). Accessed: Feb. 18, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0272884222028759>.
- [57] D. Ngo, X. Feng, Y. Huang, A. J. Rosakis, and M. A. Brown, "Thin film/substrate systems featuring arbitrary film thickness and misfit strain distributions. part i: Analysis for obtaining film stress from non-local curvature information," *International Journal of Solids and Structures*, Physics and Mechanics of Advanced Materials, vol. 44, no. 6, pp. 1745–1754, Mar. 15, 2007, ISSN: 0020-7683. DOI: [10.1016/j.ijsolstr.2006.10.016](https://doi.org/10.1016/j.ijsolstr.2006.10.016). Accessed: Nov. 25, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0020768306004215>.
- [58] M. R. Ardigo, M. Ahmed, and A. Besnard, "Stoney formula: Investigation of curvature measurements by optical profilometer," *Advanced Materials Research*, vol. 996, pp. 361–366, 2014, ISSN: 1662-8985. DOI: [10.4028/www.scientific.net/AMR.996.361](https://doi.org/10.4028/www.scientific.net/AMR.996.361). Accessed: Nov. 25, 2025. [Online]. Available: <https://www.scientific.net/AMR.996.361>.
- [59] M. Fürtsch, *Mechanical properties of thick polycrystalline silicon films suitable for surface micromachining*. Aachen: Shaker, 1999, ISBN: 978-3-8265-6594-6.
- [60] M. Xu et al., "High-strength amorphous silicon carbide for nanomechanics," *Advanced Materials*, vol. 36, no. 5, p. 10, Jan. 2, 2024, ISSN: 1521-4095. DOI: [10.1002/adma.202306513](https://doi.org/10.1002/adma.202306513). Accessed: Sep. 19, 2024. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/adma.202306513>.
- [61] K. Hajlaoui, A. B. M. Ali, M. Shabani, and S. A. Hussein, "Computational modeling of surface energy effects on linear and nonlinear frequencies in different crystalline orientations of anodic aluminum micro-beams," *Scientific Reports*, vol. 15, no. 1, p. 31 081, Aug. 24, 2025, ISSN: 2045-2322. DOI: [10.1038/s41598-025-16852-2](https://doi.org/10.1038/s41598-025-16852-2). Accessed: Nov. 21, 2025. [Online]. Available: <https://www.nature.com/articles/s41598-025-16852-2>.
- [62] P. Vashishta, R. K. Kalia, A. Nakano, and J. P. Rino, "Interaction potential for silicon carbide: A molecular dynamics study of elastic constants and vibrational density of states for crystalline and amorphous silicon carbide," *Journal of Applied Physics*, vol. 101, no. 10, p. 103 515, May 15, 2007, ISSN: 0021-8979, 1089-7550. DOI: [10.1063/1.2724570](https://doi.org/10.1063/1.2724570). Accessed: Nov. 24, 2025. [Online]. Available: <https://pubs.aip.org/jap/article/101/10/103515/923851/Interaction-potential-for-silicon-carbide-A>.
- [63] L. Hoskoti, S. S. Gupta, and M. M. Sucheendran, "Modeling of geometrical stiffening in a rotating blade—a review," *Journal of Sound and Vibration*, vol. 548, p. 117 526, Mar. 2023, ISSN: 0022460X. DOI: [10.1016/j.jsv.2022.117526](https://doi.org/10.1016/j.jsv.2022.117526). Accessed: Nov. 21, 2025. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0022460X2200709X>.
- [64] K. G. Vinod, S. Gopalakrishnan, and R. Ganguli, "Wave propagation characteristics of rotating uniform euler-bernoulli beams," *Computer Modeling in Engineering & Sciences*, vol. 16, no. 3, pp. 197–208, 2006, ISSN: 1526-1492, 1526-1506. DOI: [10.3970/cmescs.2006.016.197](https://doi.org/10.3970/cmescs.2006.016.197). Accessed: Nov. 20, 2025. [Online]. Available: <https://www.techscience-com.tudelft.idm.oclc.org/CMES/v16n3/26704>.
- [65] H. Gavin, "Geometric stiffness effects in 2d and 3d frames," Department of Civil and Environmental Engineering | Duke University, 2012, pp. 1–14. [Online]. Available: <https://people.duke.edu/~hpgavin/cee421/frame-finite-def.pdf>.