

Probing the limits of metal plasticity with molecular dynamics simulations

Luis A. Zepeda-Ruiz¹, Alexander Stukowski², Tomas Oppelstrup¹ & Vasily V. Bulatov¹

Ordinarily, the strength and plasticity properties of a metal are defined by dislocations—line defects in the crystal lattice whose motion results in material slippage along lattice planes¹. Dislocation dynamics models are usually used as mesoscale proxies for true atomistic dynamics, which are computationally expensive to perform routinely². However, atomistic simulations accurately capture every possible mechanism of material response, resolving every “jiggle and wiggle”³ of atomic motion, whereas dislocation dynamics models do not. Here we present fully dynamic atomistic simulations of bulk single-crystal plasticity in the body-centred-cubic metal tantalum. Our goal is to quantify the conditions under which the limits of dislocation-mediated plasticity are reached and to understand what happens to the metal beyond any such limit. In our simulations, the metal is compressed at ultrahigh strain rates along its [001] crystal axis under conditions of constant pressure, temperature and strain rate. To address the complexity of crystal plasticity processes on the length scales (85–340 nm) and timescales (1 ns–1 μ s) that we examine, we use recently developed methods of *in situ* computational microscopy^{4,5} to recast the enormous amount of transient trajectory data generated in our simulations into a form that can be analysed by a human. Our simulations predict that, on reaching certain limiting conditions of strain, dislocations alone can no longer relieve mechanical loads; instead, another mechanism, known as deformation twinning (the sudden re-orientation of the crystal lattice⁶), takes over as the dominant mode of dynamic response. Below this limit, the metal assumes a

strain-path-independent steady state of plastic flow in which the flow stress and the dislocation density remain constant as long as the conditions of straining thereafter remain unchanged. In this distinct state, tantalum flows like a viscous fluid while retaining its crystal lattice and remaining a strong and stiff metal.

The plasticity response of a metal depends critically on the presence or absence of dislocations before straining. Shown in Fig. 1 are the stress–strain curves computed from three different starting configurations under three different straining rates: $1.1 \times 10^7 \text{ s}^{-1}$ (our base rate, referred to as rate $\times 1$), $5.55 \times 10^7 \text{ s}^{-1}$ (rate $\times 5$) and $2.77 \times 10^8 \text{ s}^{-1}$ (rate $\times 25$); hereafter all rates are quoted in multiples of the base strain rate. A defect-free, perfect crystal yields abruptly after reaching a stress of 36–37 GPa (blue curves in Fig. 1a). Analysis of crystal configurations attained under strain reveals that, regardless of compression rate, the perfect metal yields by deformation twinning—that is, by sudden strain-induced reorientation of the crystal lattice within bounded volumes of the material⁷ (Fig. 1b). The same crystal with 24 randomly placed voids⁸ responds similarly. It yields abruptly at a lower stress of 16–19 GPa (red curves in Fig. 1a) and undergoes deformation twinning (or ‘twins’) at all strain rates that we tested (Fig. 1c). In contrast, a crystal with 24 initial dislocation loops placed into the same locations as the voids yields gradually along smooth stress–strain curves with peak stresses that are both much lower and distinctly rate-dependent (black curves in Fig. 1a). This crystal does not twin, but yields and flows exclusively through motion and multiplication of the initial dislocations, forming dense dislocation networks during

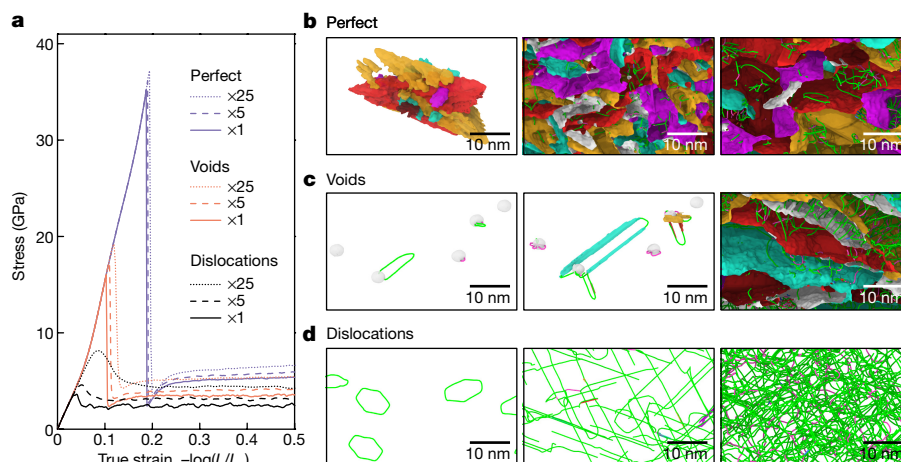


Figure 1 | Response to compression as a function of initial defect content. **a**, Stress as a function of strain. Blue, red and black curves were computed for the initially perfect crystal, the crystal with voids and the crystal with dislocations, respectively, at a temperature of 300 K. Solid, dashed and dotted lines correspond to strain rates of $\times 1$, $\times 5$ and $\times 25$, respectively. L_0 is the initial length of the sample along the z direction and L is the length at the time of the measurement. **b**, From left to right, co-nucleation of embryonic twins, twin propagation and twin growth in

the initially perfect crystal. **c**, Dislocation nucleation at voids, nucleation of twins on stretched dislocations and twin growth in the crystal with voids. **d**, Initial dislocation loops, loop extension and impingement and formation of a dense dislocation network. Dislocations appear as green and magenta lines. The twins appear as hollow volumes bounded by interfaces coloured grey, red, yellow, magenta or cyan to distinguish the parent crystal and four rotational twin variants.

¹Lawrence Livermore National Laboratory, Livermore, California, USA. ²Technische Universität Darmstadt, Darmstadt, Germany.

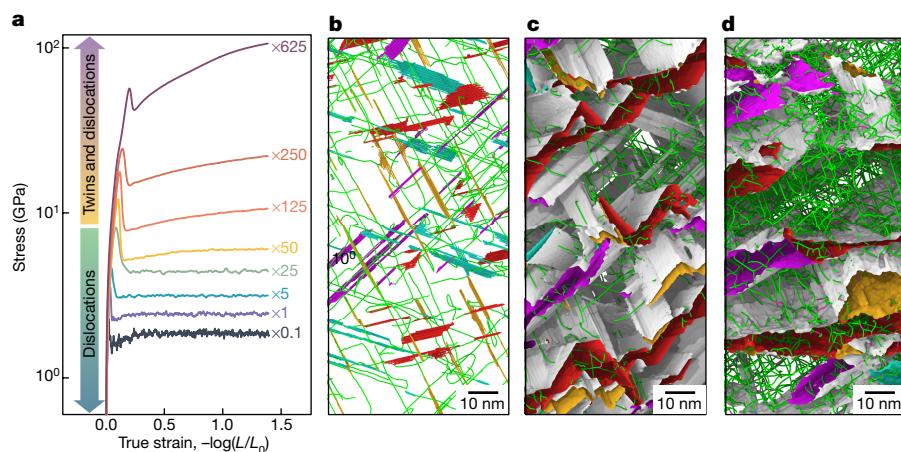


Figure 2 | Response to compression as a function of strain rate. **a**, Stress on a logarithmic scale as a function of true strain computed in molecular dynamics simulations of specimen compression at different strain rates (see colour-coded labels) and a temperature of 300 K. **b**, A snapshot of a simulation taken immediately after yield showing embryonic twins in the simulation performed at a rate of $\times 50$. **c**, A later snapshot from the same simulation showing twin propagation. **d**, An even later snapshot in which the twins have grown to span the simulation volume. Colouring of dislocation lines and twin interfaces is the same as in Fig. 1.

and soon after yielding (Fig. 1d; Supplementary Video 1). The flow response is comparable to that of ductile metals reported in laboratory experiments^{7,9}. We therefore observe that, when deprived of initial dislocations—as is common practice in atomistic simulations of metals subjected to high-rate strain—crystal strength is over-predicted by a large factor and the metal yields through mechanisms that are qualitatively different from those when dislocations are seeded in advance. Such contrasting differences are not surprising considering that multiplication of existing dislocations entails much lower stress than does nucleation of new dislocations in a perfect crystal¹⁰. Because no ductile metal has ever been obtained dislocation-free in a bulk form, in the following we confine our attention to crystals of tantalum seeded with dislocations.

Using large-scale molecular dynamics simulations, we first determine at which critical strain rate the ability of dislocations to relieve stress in tantalum becomes overwhelmed. Figure 2a shows the evolution of stress measured during compressive deformation. After an initial yield overshoot, the flow stress eventually saturates at all strain rates equal to or below $\times 25$. At any higher rate, the flow stress continuously rises through the end of compressive strain. Just like in the aforementioned perfect crystal and the crystal with voids, this strain-hardening behaviour is caused by deformation twinning (Fig. 2b–d; Supplementary Video 2). In contrast, no twinning is observed at rates of $\times 25$ and below (Fig. 1d; Supplementary Video 1). Although it is tempting to conclude that $\times 25$ is therefore the maximum rate sustainable by dislocation plasticity alone, the rate at which tantalum begins to twin varies depending on several factors, such as temperature (Fig. 3a), initial dislocation microstructure (Extended Data Fig. 1) and strain trajectory¹¹ (Methods; Extended Data Figs 2 and 3). We further observe (data not shown) that the strain rate at which twinning sets in also depends on the sign of strain (compression or tension) and the strain axis.

Although the transition to twinning is seemingly controlled by a complex interplay of the factors listed above, we observe a simple condition for the onset of twinning: it is triggered whenever the stress–strain trajectory, however complex, reaches a specific value of stress, which we refer to as a twinning threshold. The existence of a twinning threshold is in full agreement with the notion of critical resolved shear stress that is frequently encountered in the literature on crystal plasticity¹², but the applicability of which to deformation twinning in metals is still debated. In the model of tantalum examined here (Methods), twinning is observed whenever uniaxial stress reaches about 8 GPa, and this value is only weakly temperature-dependent if at all. Under [001] compression, this uniaxial stress amounts to about 1.9 GPa of shear stress resolved on eight active twinning systems of the body-centred-cubic crystal (Methods). In simulations in which strain rate or temperature are varied to gradually approach the twinning transition, dislocations have ample time to multiply and their density reaches around 10^{17} m^{-2} . This limiting value closely matches a prediction given for the maximum

dislocation density a crystal can sustain before collapsing¹³ (see Methods and Extended Data Fig. 4 for more details on this dislocation density limit).

What happens once twinning is triggered depends on the initial material state¹⁴ and on the strain trajectory, both before and after twinning. For example, at a rate of $\times 25$ and temperature of $T = 300 \text{ K}$, and at a rate of $\times 5$ and $T = 50 \text{ K}$, we observe that twinning is triggered momentarily when the stress–strain trajectory first crosses 8 GPa, but the newly nucleated (or ‘embryonic’) twins subsequently recede once the flow stress drops below this threshold value (Figs 1d and 3c). Very different defect microstructures develop at higher rates at $T = 300 \text{ K}$ (Fig. 2d) or at lower temperatures at a rate of $\times 5$ (Fig. 3b) at which the twins grow rapidly, filling large fractions of material volume, impinging on each other and even suffering secondary twinning (twins inside twins). Yet another distinct scenario is observed in simulations in which the twinning threshold is reached by slowly increasing (ramping) the strain rate at a constant temperature or, alternatively, by continuously cooling the metal while it is strained at a constant rate. Once the stress reaches about 8 GPa, numerous twins nucleate but do not grow and remain embryonic, presumably constrained from further growth by the dense dislocation networks that are established by then (see Methods and Extended Data Fig. 3 for more detail on simulations under continuous cooling). The fact that the threshold stress is nearly

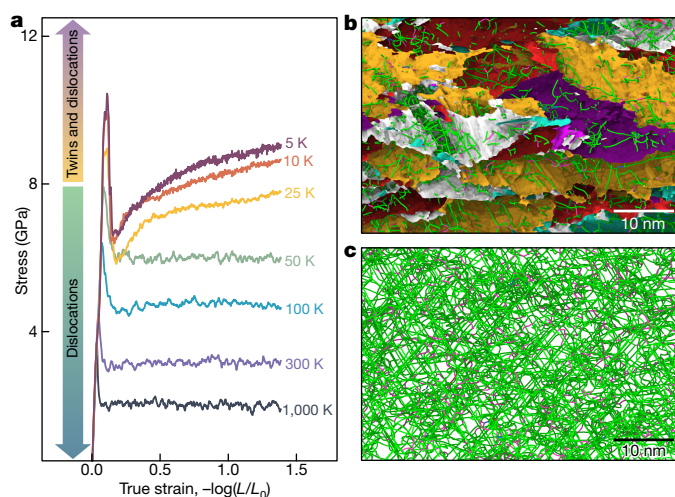


Figure 3 | Response to compression as a function of temperature. **a**, Stress as a function of true strain computed in molecular dynamics simulations of specimen compression at different temperatures (see colour-coded labels) at a rate of $\times 5$. **b**, A snapshot extracted from a simulation at $T = 25 \text{ K}$ depicting well-developed twins surrounded by dislocations. **c**, A snapshot extracted from a simulation at $T = 50 \text{ K}$ showing a dense dislocation network but no twins. Colouring of lines and interfaces is the same as in Fig. 1.

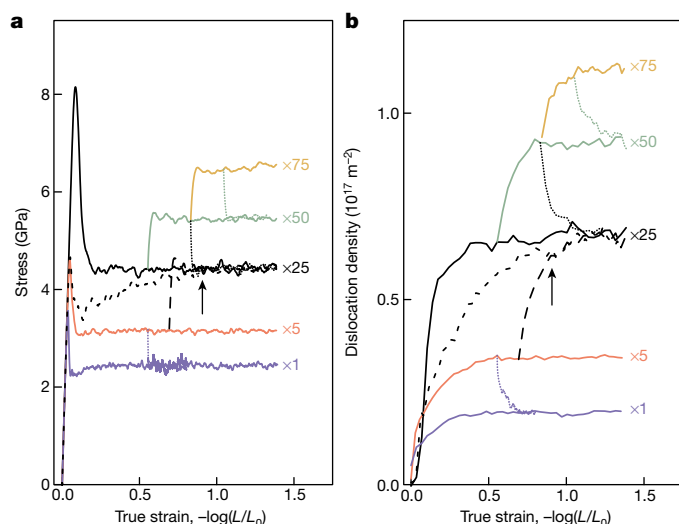


Figure 4 | Attaining a path-independent flow state. **a**, Stress–strain response computed in molecular dynamics simulations along four different types of strain trajectory at $T = 300$ K: solid lines correspond to compression at constant rates, dotted lines show stress relaxations after jumps to lower rates, the long-dashed line shows stress relaxation after a jump to a higher rate and the short-dashed line corresponds to a strain rate increasing continuously from zero to $\times 25$ and then being held constant from the strain marked by the vertical arrow. **b**, Dislocation density as a function of strain computed along the same strain trajectories. Line colours and types match the stress–strain trajectories depicted in **a**.

independent of the strain conditions is explained by our observation that twinning nucleation always occurs by stress-driven extension of screw dislocations, as has previously been suggested^{15,16}. However, in the presence of other possible pathways for the initiation of twinning, such as grain boundaries or triple lines in polycrystals, the twinning threshold may have a different value or may even be impossible to define.

We now turn to our simulated strain trajectories in which stress stays below the threshold and twinning is not triggered. We find that any such trajectory leads to an ultimate steady state of material flow that is then maintained indefinitely as long as the strain conditions (strain rate, temperature and strain axis) remain unchanged. In Fig. 4 we present substantial evidence that, no matter how the final strain conditions are attained—by straining at a constant rate throughout, in a sequence of rate jumps or under a continuously varying rate—the state that the material ultimately attains is manifestly path-independent. For example, the four different types of strain trajectory that lead to the same limiting rate of $\times 25$ clearly converge to the same stress (black lines in Fig. 4a) and dislocation density (black lines in Fig. 4b), the values of which depend on only the final strain rate. We further show in a ‘metal kneading’ simulation (Methods; Extended Data Fig. 5; Supplementary Video 3) and other analyses that flow stress, dislocation density and other more detailed quantitative characteristics of dislocation microstructures remain constant over arbitrarily large compressive strains (Methods; Extended Data Figs 6–8). We therefore conclude that the ultimate ‘state of flow’ that is predicted remains steady as long as strain conditions remain unchanged.

We regard the flowing metal as an open thermodynamic system maintained at a constant rate of energy dissipation in which new mechanical energy is supplied by strain, then converted through dislocation multiplication, annihilation and motion into heat and eventually collected and removed by the thermostat. In this steady state, dislocation multiplication is exactly balanced by dislocation annihilation. A thermodynamic potential and an equation of state¹⁷ should therefore exist that uniquely relate the internal state variables, such as saturated flow stress and saturated dislocation density, to the external strain conditions, that is, temperature, pressure, strain rate and axis. Shown in Fig. 5a are two such equations of flow state, for the flow stress and

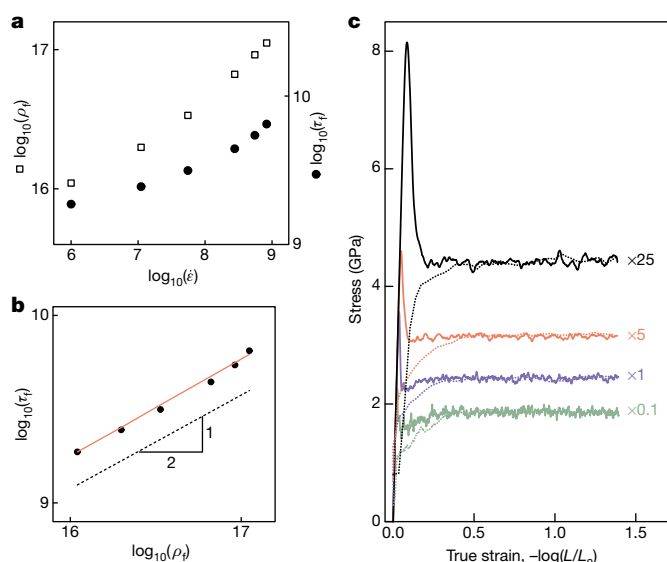


Figure 5 | Flow stress and dislocation density at $T = 300$ K. **a**, Saturated flow stress τ_f (filled circles) and saturated flow density ρ_f (open squares) as functions of strain rate $\dot{\epsilon}$. **b**, τ_f versus ρ_f attained at different rates (filled circles). The black dashed line is the Taylor equation with coefficient $\alpha = 0.3$ taken from experiment⁹; this line has a slope of $1/2$ (as indicated). The solid red line is a Taylor equation corrected for the presence of strong ternary junctions, the fraction of which increases with increasing strain rate (Methods). **c**, Stress–strain response computed in molecular dynamics simulations at different rates (solid lines; see colour-coded labels), and stress predicted from the concomitant dislocation density using the Taylor equation along the same strain trajectories (dotted lines).

the flow density, computed in our molecular dynamics simulations as functions of strain rate while holding all other control parameters fixed.

Straining rates and dislocation densities encountered in our simulations are common to high-rate compression experiments^{18,19}. However, direct comparisons to such dynamic tests are uncertain owing to their complex and largely unknown strain trajectories that are far from stationary^{20,21} and only indirectly inferred. Here, we opted for fully controlled simulations in which intrinsic plasticity properties are computed directly, in the most straightforward manner possible, akin to the standard low-rate (10^{-5} s^{-1} to 10^2 s^{-1}) laboratory tests of crystal plasticity. It remains to be shown whether it is sensible to compare our predictions to low-rate experiments, given the orders of magnitude difference in the strain rates. For instance, although the notion of saturated flow is not entirely new^{22–24}, it is unclear how our prediction of a steady flow state matches up with the expectation of ‘hardening forever’ that prevails in the low-rate literature. Few low-rate compression experiments reach strains of order unity, which are necessary to observe saturation, and those that do are commonly affected by extrinsic factors that conceal the true response, including barrelling or necking and rotation of the strain axis^{25–27}. However, in at least one experiment, in which strain conditions remained unchanged over very large strains²⁸, saturated flow at a manifestly constant stress has been observed in single-crystal molybdenum (a close cousin of tantalum) compressed along the same [001] axis at a rate of 10^{-3} s^{-1} . We therefore regard our prediction of a saturated flow state as a tantalizing, general hypothesis that merits further scrutiny.

The above comparison draws attention to a more general question: whether the physics of metal plasticity at high rates differs from the much studied mechanisms that define plasticity under quasi-static strain conditions. We do observe certain differences. As shown in Fig. 5a, our predicted saturated flow stress and saturated dislocation density are markedly more rate-dependent than the low-rate experimental data for tantalum would suggest⁷. Another notable difference concerns the Taylor hardening equation that relates flow stress τ to dislocation density ρ as $\tau = \alpha \mu b \sqrt{\rho}$, where μ and b are the shear

modulus and the magnitude of the Burgers vector, respectively, and α is a dimensionless constant believed to reflect the strength of dislocation interactions²⁹. Low-rate straining experiments⁷ on single-crystal tantalum show that the Taylor equation holds reasonably well at $\alpha = 0.3$. However, extrapolated to the considerably higher values of saturated flow density ρ_f that are observed in our simulations, the low-rate empirical equation predicts a saturated flow stress τ_f that is markedly lower than our simulation results (Fig. 5b). To approximately fit our molecular dynamics simulation data, the Taylor coefficient must be higher, $\alpha = 0.45$. Furthermore, as can be seen in Fig. 5b, at the highest rates the flow stress deviates upward from the standard square-root prediction (black dashed line with a slope of 1/2 in the logarithmic plot). We relate these discrepancies to the recently discovered ternary junctions that tie together three dislocations and thus are much stronger than the ordinary binary junctions that hold together only two dislocations³⁰. As detailed in Methods, the fraction of strong ternary junctions is observed to grow markedly with increasing strain rate, which may explain both the increased hardening coefficient α and the upward deviation from the standard Taylor relationship.

In our simulations, deviations from Taylor hardening are even more pronounced during the transient strain regime that leads to the steady flow state (Fig. 5c). Here, the flow stress clearly outpaces the density evolution to the extent that at the highest strain rate of $\times 25$ the flow stress shows no hardening at all and appears to reach saturation much sooner than does the dislocation density. With decreasing strain rates, however, the flow stress and the flow density appear to increase and saturate more and more in lock step. This trend points to an increasingly higher resistance to dislocation motion with increasing strain rates, owing to two sources: greater dislocation–lattice coupling and increased background (or ‘forest’) resistance resulting from the increased frequency of dislocation collisions. At high rates, lattice resistance dominates in the early stages of straining when dislocations are relatively scarce (and dislocation collisions are infrequent), and forest resistance catches up, but with a notable delay, as dislocations multiply.

Apart from the aforementioned quantitative differences, in our molecular dynamics simulations the dislocations behave as expected at low strain rates: (1) they move in a stop-and-go fashion, remaining relatively motionless against obstacles (junctions and self-pinning cusps³¹) for periods of time with intermittent bursts of faster motion; (2) the distribution of the orientation of dislocation lines is non-uniform and dominated by screw lines with orientations nearly parallel to their Burgers vectors (see Methods); and (3) the average dislocation velocities range between 1 m s^{-1} and 60 m s^{-1} , placing the dominant screw dislocations into the same temperature- and stress-activated motion regime as expected at low rates³². We thus observe that, quantitative differences notwithstanding, the basic mechanisms of dislocation-mediated plasticity in tantalum seem to remain the same for strain rates from the quasi-static to dynamic shock regime, that is, over the full range of 14 decades of experimentally accessible strain rates. If this proves to be the case, then simulations of the kind reported here present a distinct and exciting opportunity to examine and understand the fundamental atomistic mechanisms that underlie metal plasticity in unprecedented detail (Extended Data Fig. 9; Supplementary Video 4).

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 28 April 2016; accepted 23 June 2017.

Published online 27 September 2017.

1. Bulatov, V. V. & Cai, W. *Computer Simulations of Dislocations* 196–240 (Oxford Univ. Press, 2006).
2. Ghoniem, N. M. *et al.* Parametric dislocation dynamics: a thermodynamics-based approach to investigations of mesoscopic plastic deformation. *Phys. Rev. B* **61**, 913–927 (2000).
3. Feynman, R. *Lectures on Physics* Vol. 1, 3–6 (1963).
4. Stukowski, A. Visualization and analysis of atomistic simulation data with OVITO — the Open Visualization Tool. *Model. Simul. Mater. Sci. Eng.* **18**, 015012 (2010).

5. Stukowski, A. & Albe, K. Extracting dislocations and non-dislocation crystal defects from atomistic simulation data. *Model. Simul. Mater. Sci. Eng.* **18**, 085001 (2010).
6. Christian, J. W. & Mahajan, S. Deformation twinning. *Prog. Mater. Sci.* **39**, 1–157 (1995).
7. Hoge, K. G. & Mukherjee, A. K. The temperature and strain rate dependence of the flow stress of tantalum. *J. Mater. Sci.* **12**, 1666–1672 (1977).
8. Tramontina, D. *et al.* Molecular dynamics simulations of shock-induced plasticity in tantalum. *High Energy Density Phys.* **10**, 9–15 (2014).
9. Mitchell, T. E. & Spitzig, W. A. Three-stage hardening in tantalum single crystals. *Acta Metall.* **13**, 1169–1179 (1965).
10. Frank, F. C. & Read, W. T. Jr. Multiplication processes for slow moving dislocations. *Phys. Rev.* **79**, 722–723 (1950).
11. Meyers, M. A. *et al.* The onset of twinning in metals: a constitutive description. *Acta Mater.* **49**, 4025–4039 (2001).
12. Shields, J. A. *et al.* Deformation of high purity tantalum single crystals at 4.2 K. *Mater. Sci. Eng.* **20**, 71–81 (1975).
13. Cotterill, R. M. J. Does dislocation density have a natural limit? *Phys. Lett. A* **60**, 61–62 (1977).
14. Florando, J. N. *et al.* Effect of strain rate and dislocation density on the twinning behavior in tantalum. *AIP Adv.* **6**, 045120 (2016).
15. Sleswyk, A. W. 1/2(111) screw dislocations and nucleation of {112}{111} twins in the b.c.c. lattice. *Phil. Mag.* **8**, 1467–1486 (1963).
16. Marian, J., Cai, W. & Bulatov, V. V. Dynamic transitions from smooth to rough to twinning in dislocation motion. *Nat. Mater.* **3**, 158–163 (2004).
17. Martyushev, L. M. & Seleznev, V. D. Maximum entropy production principle in physics, chemistry and biology. *Phys. Rep.* **426**, 1–45 (2006).
18. Hsiung, L. L. Shock-induced phase transformation in tantalum. *J. Phys. Condens. Matter* **22**, 385702 (2010).
19. Nemat-Nasser, S. *et al.* Microstructure of high-strain, high-strain-rate deformed tantalum. *Acta Mater.* **46**, 1307–1325 (1998).
20. Lu, C. H. *et al.* Phase transformation in tantalum under extreme laser deformation. *Sci. Rep.* **5**, 15064 (2015).
21. Reed, B. W. *et al.* A unified approach for extracting strength information from non-simple compression waves. Part II. Experiment and comparison with simulation. *J. Appl. Phys.* **110**, 113506 (2011).
22. Johnston, W. G. & Gilman, J. J. Dislocation velocities, dislocation densities and plastic flow in lithium fluoride crystals. *J. Appl. Phys.* **30**, 129–144 (1959).
23. Mecking, H. & Kocks, U. F. Kinetics of flow and strain hardening. *Acta Metall.* **29**, 1865–1875 (1981).
24. Seeger, A. Evidence of enhanced self-organization in the work-hardening stage V of fcc metals. *Philos. Mag. Lett.* **81**, 129–136 (2001).
25. Banerjee, J. K. Barreling of solid cylinders under axial compression. *J. Eng. Mater. Technol.* **107**, 138–144 (1985).
26. Roylance, D. *Stress–Strain Curves* https://ocw.mit.edu/courses/materials-science-and-engineering/3-11-mechanics-of-materials-fall-1999/modules/MIT3_11F99_ss.pdf (MIT OpenCourseWare, 2001).
27. Khan, A. S. & Huang, S. *Continuum Theory of Plasticity* 37–40 (Wiley-Interscience, 1995).
28. Carpay, F. M. A. *et al.* Constrained deformation of molybdenum single crystals. *Acta Metall.* **23**, 1473–1478 (1975).
29. Saada, G. On hardening due to the recombination of dislocations. *Acta Metall.* **8**, 841–847 (1960).
30. Bulatov, V. V. *et al.* Dislocation multi-junctions and strain hardening. *Nature* **440**, 1174–1178 (2006).
31. Louchet, F. & Viguier, B. Ordinary dislocations in γ -TiAl: cusp unzipping, jog dragging and stress anomaly. *Philos. Mag. A* **80**, 765–779 (2000).
32. Christian, J. W. & Masters, B. C. Low-temperature deformation of body-centered cubic metals. *Proc. R. Soc. Lond. A* **281**, 223–239 (1964).

Supplementary Information is available in the online version of the paper.

Acknowledgements This work was performed under the auspices of the US Department of Energy by Lawrence Livermore National Laboratory under contract W-7405-Eng-48. This work was supported by the NNSA ASC programme. Computing support for this work came from the Lawrence Livermore National Laboratory (LLNL) Institutional Computing Grand Challenge programme and Jülich Supercomputing Centre at Forschungszentrum Jülich, Germany.

Author Contributions L.A.Z.-R. ran most of the molecular dynamics simulations and analysed the results, A.S. ran molecular dynamics simulations and developed methods for computational microscopy and visualization, T.O. optimized run-time efficiency and data management of molecular dynamics simulations, and V.V.B. developed the concept, planned the research, generated starting configurations for molecular dynamics simulations, analysed the results and wrote the paper.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations. Correspondence and requests for materials should be addressed to V.V.B. (bulatov1@llnl.gov).

Reviewer Information Nature thanks M. Zaiser and the other anonymous reviewer(s) for their contribution to the peer review of this work.

METHODS

Molecular dynamics (MD) simulations. The key challenge for MD simulations of phenomena as complex as crystal plasticity is their high computational costs, which place severe limits on accessible length scales and timescales. Here we utilize some of the world's leading HPC facilities to push beyond these scale limits. Simulations were performed using the open-source code LAMMPS³³ on brick-shaped fragments of tantalum single crystals seamlessly embedded into an infinite crystal under three-dimensional periodic boundary conditions. Interaction among tantalum atoms was modelled with the embedded atom method (EAM) interatomic potential³⁴. The initial aspect ratio of the brick-shaped simulation volume was 1:2:4, with dimensions ranging from $L_x \times L_y \times L_z = 128a_0 \times 256a_0 \times 512a_0$ (about 33 million atoms) to $L_x \times L_y \times L_z = 256a_0 \times 512a_0 \times 1,024a_0$ (about 268 million atoms), oriented along the principal [100], [010], [001] axes of the body-centred-cubic lattice, respectively ($a_0 = 0.33$ nm is the lattice constant of tantalum at ambient conditions). The crystal was compressed either at a constant rate or at a variable 'true' rate

$$\dot{\epsilon}_T = \frac{1}{L_z} \frac{dL_z}{dt}$$

along its initially longest [001] axis L_z , while its two lateral dimensions L_x and L_y were left to expand to maintain the 'xx' and 'yy' components of the mechanical stress near zero. Most straining simulations were continued until the 'longest' dimension L_z was compressed to one-quarter of its initial size, by which time L_x and L_y had almost doubled, thus keeping volume and pressure nearly constant throughout the simulation. Temperature was maintained constant using a Langevin thermostat with the damping parameter set to 10 ps. Everywhere in the text by 'stress' we mean the uniaxial compressive stress σ_{zz} that develops in the model crystal in response to strain along its [001] z axis.

In situ computational microscopy. MD simulations of crystal plasticity generate enormous amounts of data: a simulation of the kind reported here generates about 5 exabytes (1 exabyte = 10^{18} bytes) of atomic trajectory data in just one day on the Sequoia supercomputer³⁵, an amount comparable to the estimated worldwide storage capacity of Google. To be useable, MD data must be compressed to reduce its size and to recast it into a form that a human can grasp and comprehend. Here we rely on recently developed methods of *in situ* computational microscopy to identify and precisely characterize extended defects in crystals^{4,5} and to reduce the amount of data by orders of magnitude. In addition to the standard local atom filtering techniques, we use the dislocation extraction algorithm (DXA)⁵ and the grain segmentation algorithm (GSA) to reveal structural defects in simulated crystals and to reduce the amount of data to deal with by orders of magnitude. The DXA accurately traces and indexes the dislocation lines and builds a concise network representation of the crystal microstructure identical to the one used in the mesoscale 'discrete dislocation dynamics' method. The newly developed GSA automatically partitions the atomistic crystal into regions of similar crystal orientation (grains) and generates a geometric representation of the interfaces delineating the grains, thus identifying any twinning. This conversion from the fully atomistic model to a highly reduced description of the defect microstructure is performed at regular time intervals during the MD simulation and enables us to follow the motion and reactions of dislocations, the onset of twinning and other important events with an arbitrary degree of detail. Defect microstructures shown in Figs 1–3, Extended Data Fig. 9 and Supplementary Videos 1–4 were rendered with the program OVITO⁴.

As an example, Extended Data Fig. 9 depicts several important events in the life of dislocations detected and rendered using the 'atomic strain' calculation function available in OVITO. The atomic strain function requires the user to supply two atomistic configurations, initial and final, from the same simulated trajectory. When invoked, this function compares positions of atoms in the two supplied configurations and computes a local per-atom measure of the von Mises shear strain for every atom. To see where atomic scale slip has taken place, the user can select for subsequent rendering only the atoms whose shear strain is above a user-defined threshold and delete all other atoms for clarity. To generate the snapshots shown in Extended Data Fig. 9, the 'high-strain' atoms are juxtaposed with two dislocation networks extracted from the same initial and final configurations using the DXA analysis tool in OVITO (DXA requires only one atomic configuration). The snapshots are extracted from the MD simulation shown in Supplementary Video 4.

Twin nucleation events. Despite similar eventual outcomes—twinning—the perfect, defect-free crystal (crystal 'A') and the crystal with voids (crystal 'B') respond to identical strain conditions differently. Whereas crystal A twins at a peak stress that is nearly rate-independent (36–37 GPa), crystal B twins at a lower peak stress that depends more noticeably on the strain rate. Twinning initiation in

crystal A bears signatures of lattice instability: twins are nucleated homogeneously, springing up simultaneously and randomly anywhere in the crystal volume. The small, fish-shaped twin embryo seen in the leftmost frame in Fig. 1b consists of sub-particles of four different colours: all four rotational twin variants co-nucleate simultaneously, perhaps owing to the high symmetry of the [001] strain axis. Subsequently, numerous twins rapidly fill much of the crystal volume (middle frame in the same sequence) and grow seemingly larger in the later stages of straining (right frame). In contrast, twinning in crystal B is a two-stage process. First, dislocations are emitted from the pre-existing voids³⁶ (leftmost frame in Fig. 1c); and second, twins heterogeneously (at preferred positions) nucleate on the emitted dislocations according to Sleeswyk's mechanism^{15,16} (middle frame of the same sequence). Moderate rate dependence of the yield stress in the crystal with voids reflects an activated time-dependent nature of dislocation nucleation from the voids. Following initial emission of dislocation loops from the voids, formation of embryonic twins does not appear to be temperature- or stress-activated, but occurs as soon as the loops draw sufficiently long screw segments.

Yield behaviour depends on initial dislocation sources. In Extended Data Fig. 1 we show stress–strain curves computed in three MD simulations of compressive straining at the same rate $\times 25$, but starting from three different initial configurations of dislocation sources (loops). The blue curve is the response of a crystal with four hexagon-shaped loops randomly placed in the simulation volume, one loop for each of the four different Burgers vectors of the $1/2\langle 111 \rangle$ type. The loops are of the vacancy type and approximately 8 nm in diameter corresponding to an initial dislocation density of about $\rho_0 = 1.5 \times 10^{14} \text{ m}^{-2}$. The red curve is the response of a crystal with 24 loops of the same type and size randomly placed in the volume, six loops per Burgers vector ($\rho_0 = 9 \times 10^{14} \text{ m}^{-2}$). The black curve is the response of a crystal with four initial loops of six times larger diameter, one loop per Burgers vector ($\rho_0 = 9 \times 10^{14} \text{ m}^{-2}$). Even though there is some dependence on the loop size (compare red and black curves), yield stress is far more sensitive to the initial dislocation density (compare blue and black curves). Yielding of the crystal with the lower initial dislocation density entails nucleation and growth of twins concomitant with strain hardening, whereas two other crystals of the higher initial dislocation density yield by dislocation multiplication and motion alone, produce no twins and exhibit no strain hardening. Our simulations suggest the following inter-relationship between the initial source density and the yield behaviour: the lower the initial density the longer time it takes for the dislocations to multiply and to prevent twinning¹⁴. Conversely, in the unrealistic hypothetical limit of zero initial dislocation, density twinning is perhaps unavoidable (see Fig. 1b, c).

Detection of twinning transition. As can be seen from Fig. 2a, at temperature 300 K the crystal does not twin when compressed at rate $\times 5$. However, the plasticity response of tantalum depends sensitively on temperature as illustrated by the stress–strain curves shown in Fig. 3a that were obtained in MD simulations of compressive straining at the same rate $\times 5$ but at temperatures ranging from 5 K to 1,000 K. In addition to a marked increase in the yield and flow stress with decreasing temperature, the response curves at 25 K and still lower temperatures are qualitatively different from the stress–strain curves obtained at higher temperatures: here again, hardening observed past yield signals twinning. Using the GSA we confirmed that under compressive strain at rate $\times 5$ the crystal does indeed twin at 25 K (Extended Data Fig. 2), and at 10 K and 5 K. This observation is in qualitative agreement with experimental data showing that strain rates at which twinning is observed increase with increasing temperature. We therefore expect, with all other conditions fixed, that a separatrix exists in the rate–temperature plane delineating two essentially different modes of plastic response: 'dislocations only' and 'dislocations + twins'.

It is possible to establish where the separatrix is by running a series of MD simulations at a fixed temperature and many different rates, as we did to obtain the data shown in Fig. 2a. Alternatively, the same twinning transition can be searched for at a fixed rate, by running simulations at many different temperatures, as we did to obtain data for Fig. 3a. However, both approaches are computationally expensive and may not be sufficiently accurate: the accuracy of pinning down the twinning transition depends on the spacing of the rate–temperature grid and it is not necessarily known in advance where in the rate–temperature plane to focus the computational effort. Furthermore, occurrence of twinning depends on various extrinsic factors such as variations in the initial dislocation microstructure and on a multitude of possible time-dependent strain trajectories. To avoid such uncertainties, we identify an intrinsic twinning transition with the minimal rate at which twinning is triggered in the limit of adiabatically slow rate ramping at a constant temperature. Conversely, the same intrinsic transition can be defined as the minimal temperature at which twinning is not observed in the limit of adiabatically slow cooling under a constant strain rate. Although it is not practically possible to perform a simulation (or experiment) over an infinitely long time, the above definitions suggest two efficient protocols for detecting the twinning

transition. It can be more computationally efficient and/or accurate to detect and locate a twinning transition by: (1) raising the strain rate as slowly as possible at a fixed temperature or (2) cooling the crystal as slowly as possible at a fixed strain rate. In both methods, rate or temperature ramping should start on the ‘dislocations only’ side of the separatrix. Obviously, the better one knows (or anticipates to know) where to expect the transition, the more accurately the transition itself will be located.

We tried both simulation protocols to detect the twinning transition temperature at a fixed rate and to detect the twinning transition rate at a fixed temperature. An example is presented in Extended Data Fig. 3, in which the twinning fraction—the fraction of atoms inside all twin particles computed using the GSA—is plotted as a function of temperature. A robust increase in the volume fraction of twins observed at temperatures of about 30–50 K signals the onset of twinning. The twinning transition temperature estimated from this simulation is consistent with the data presented in Fig. 3a, in which twinning was observed at 25 K but not at 50 K. Furthermore, continued cooling causes the flow stress to increase and to reach the familiar threshold value of about 8 GPa in the same range of temperatures in which the twin fraction begins to increase (compare to Fig. 2a). A continuous increase in the strain rate at a fixed temperature is also observed to result in twinning; however, the twin fraction remains small, which makes detection and quantification of the twinning transition less robust. Visual analyses reveal that, under both continuous cooling and continuous rate ramping, twins nucleate once a sufficiently high stress of about 8 GPa is reached, but do not grow and remain embryonic even as the stress continues to increase. We leave it to future analysis to establish exactly why the twins remain undersized, but our current hypothesis is that twin growth is somehow arrested by very dense dislocation networks that develop in simulations under continuously increasing strain rates or under continuous cooling.

Taken together, the simulation data presented here and in the main text (Figs 2 and 3) indicate that, no matter how the twinning transition is induced—under constant rate and temperature, by rate ramping or by continuous cooling—the stress at which twins first appear is only weakly dependent on temperature and rate. In fact, to zeroth order, stress to trigger twinning can be taken as approximately constant, about 8 GPa, within the range of strain rates and temperatures that we explored in our simulations. Given the Schmid factor of eight twinning systems active under [001] compression, $1/\sqrt{18} = 0.2357$, the resolved (shear) threshold stress that triggers twinning in our model of tantalum is estimated to be 1.9 ± 0.3 GPa. This parameter can be regarded as a material constant—the critical resolved shear stress for twinning¹²—in conditions when twinning is initiated on pre-existing screw dislocations via the mechanism suggested in refs 15 and 16 as observed in our simulations.

Post-straining reduction of dislocation density. Under the highest strain rates explored in our simulations, the dislocation density reaches extreme levels of about 10^{17} m^{-2} consistent with an earlier estimate for the limiting density¹³. Although dislocation densities of such magnitude have been reported, invariably, TEM experiments reveal only *ex situ* (post-mortem) density whereas our data pertains to *in situ* (during straining) density. Owing to several mechanisms of density relaxation (recovery) coming into play after straining, the *ex situ* density can be much lower than the *in situ* density. In high-rate straining experiments in which *in situ* dislocation density reaches extreme values, post-mortem density relaxation can be particularly large and should not be ignored. Here we present computational experiments intended to illustrate, within our limited computational resources, two factors known to contribute to post-straining density reduction.

The strain rates explored in our simulations are common to high-rate compression experiments; however, strain trajectories encountered in such dynamic tests are time-dependent, highly convoluted or even unknown, making direct comparisons with our fully controlled MD simulations impossible at present. What is known, however, is that fast multiplication and motion of dislocations under high stress generates much heat that has little or no time to disperse during straining, thus greatly affecting the material response^{20,21}. Owing to computational limitations, it is impossible at present to simulate such dynamical tests on their realistic timescales and length scales, but we can try to gauge the possible effects of rapid heat production. Extended Data Fig. 4 shows the kinetics of density relaxation following two different strain trajectories, both ending at the same rate of $\times 75$. The upper curve represents kinetics of density relaxation following isothermal straining at rate $\times 75$, just below the twinning transition at $T = 300$ K. It is in this simulation that the dislocation density reaches its highest level of $1.1 \times 10^{17} \text{ m}^{-2}$ while still not triggering twinning. The lower curve represents density relaxation following adiabatic compression at a rate linearly increasing from 0 to the same end rate of $\times 75$. In the isothermal relaxation simulation, temperature was maintained close to 300 K, whereas in the adiabatic simulation heat generated during plastic straining was retained in the system resulting in temperature rising to 1,350 K.

Qualitatively, the two kinetics can be expected to bound the (unknown) material response to straining at rate $\times 75$ because experimental high-rate tests are neither fully isothermal nor fully adiabatic.

As soon as straining is stopped, stress is observed to relax to zero over just a few tens of picoseconds in both isothermal and adiabatic simulations. Over the same short time interval, the dislocation density rapidly decreases owing largely to fast straightening (springing back or recoiling) of dislocations that were previously bowed out under stress. Following this initial reduction, more extended (lazy) relaxation kinetics commences, manifesting itself in a continuing downward slope of dislocation density plotted as a function of the logarithm of time. Assuming continued relaxation to follow the same lazy, stretched exponential kinetics all the way to its infinite-time asymptotic value of zero³⁷, the dislocation density can be expected to decrease by an additional order of magnitude on the macroscopic timescales of seconds and days.

Still another factor that can reduce dislocation densities observed in experiment is prepping strained specimens for subsequent TEM observations. Here we emulate the effect of specimen thinning by interrupting the isothermal density relaxation simulation, replacing the periodic boundary conditions along the x axis with free-surface boundary conditions and continuing isothermal density relaxation, but now with two opposite faces of the simulation volume exposed to vacuum. At the time of the switch, the dimension of the crystal along its x axis is 83 nm, which is appropriately close to a typical thickness of an electron-transparent film used in TEM experiments. As seen from the plots, reduction in the dislocation density during post-straining relaxation in the ‘thin film’ is noticeably greater than in the ‘bulk’ 3D periodic crystal.

Computational experiments discussed in this section, however simple and limited in their time horizon, generally substantiate the expectation that volume-averaged dislocation densities observed post-mortem in crystals subjected to extreme-rate compression should be substantially lower—perhaps by an order of magnitude if not more—than the *in situ* densities.

Additional simulations and analyses confirm that the flow state is stationary.

In the MD simulation that we refer to as ‘metal kneading’, a brick-shaped volume of tantalum crystal with initial dimensions $x:y:z = 1:2:4$ is kneaded like a piece of dough. First, the brick is compressed along its longest, z dimension until its dimensions become 2:4:1 (owing to expansion in its two lateral dimensions, the volume of the brick remains nearly constant under compression, reflecting the volume-preserving nature of metal plasticity). Then, another simulation starts in which the same brick is compressed along its now longest y axis so that its final dimensions become 4:1:2. After completing the second compression cycle, the brick is compressed along its now longest x axis, and so on. Stress–strain curves shown in Extended Data Fig. 5 demonstrate that flow stress remains the same over four such consecutive compression cycles (see Supplementary Video 3). Although not shown in the plot, the dislocation density is also observed to remain unchanged along the same strain trajectory, except for short transients in the beginning of each compression cycle. We continued the metal kneading process well beyond the initial four cycles depicted in Extended Data Fig. 5 and observed no changes in the flow behaviour.

To further confirm that the saturated flow observed in our MD simulations is indeed stationary, we examine evolution under strain of other, more detailed characteristics of material state, namely dislocation network topology and dislocation line geometry.

Aside from occasional disconnected loops, dislocations in a crystal form networks in which junction nodes join together three or more dislocations. In close analogy with electric circuits, dislocation networks conserve vector current: (a) the Burgers vector is conserved along each dislocation line and (b) the Burgers vectors of dislocations joined at each node sum to zero. As a rule, bulk metals already contain dislocation networks in their as-grown state. When a crystal is strained, dislocations multiply and collide increasingly more frequently, resulting in a gradual increase in the dislocation network density; this is precisely what we observe in our MD simulations. It has been tacitly assumed over the years that every dislocation network is built from the same basic unit—a junction in which two parent dislocations merge together to form a third (product) dislocation. However, the recent discovery of dislocation multi-junctions has led to reconsideration of this basic tenet, bringing to the fore a new object—a ternary junction in which three dislocations merge together to form a fourth dislocation. The importance of ternary junctions is in their strength: ternary junctions hold dislocations together much more strongly than do ordinary (binary) junctions³⁰. If present in substantial concentrations, ternary reactions should add considerably to the overall network resistance to dislocation motion and thus raise the flow stress.

To track evolution of dislocation networks along our simulated strain trajectories, we developed an algorithm consisting of a sequence of unit topological operations that deconstruct an arbitrary dislocation network into its constituent

binary and ternary junctions. In Extended Data Fig. 6 we show the numbers of (ordinary) binary junctions and ternary junctions computed for a sequence of network configurations observed along a single strain trajectory at our base rate $\times 1$ ($1.11 \times 10^7 \text{ s}^{-1}$). As expected, initially there are no junctions of either type because our initial configuration contains only isolated dislocation loops. Then, as soon as dislocations multiply sufficiently to start intersecting each other, the number of binary junctions begins to increase rapidly, followed closely by an increase in the number of ternary junctions. Further confirming a stationary nature of the flow state, both densities appear to saturate at a strain of about 0.4. Two aspects of network formation kinetics are noteworthy: (1) the initial increase in the density of ternary junctions trails closely behind the kinetics of ordinary (binary) junctions, which makes sense because the ternary junctions are expected to result from collisions of third dislocations with binary junctions that are already present; and (2) in the stationary (saturated) flow state, the density of ternary junctions becomes as high as that of the ordinary junctions. The latter observation is particularly startling and suggests that the overall strength of a fully formed (saturated) dislocation network may be largely defined by the ternary junctions, owing to their greater strength.

It is of interest to see if and how network composition in the saturated state depends on the strain rate. In Extended Data Fig. 7 we show the numbers of binary and ternary junctions found within the dislocation network evolving along a strain trajectory that begins at strain rate $\times 25$ and then on reaching strain about 0.55 jumps to rate $\times 50$, followed by another jump to rate $\times 75$ at strain about 0.85. Clearly, the numbers of junctions of each type (binary and ternary) closely follow the strain trajectory even if it takes some time (and strain) for the network composition to settle after each rate jump. It is equally clear that the fraction of ternary (multi-) junctions increases with the increasing strain rate, reaching 57% at rate $\times 75$, which, as we observed, is the highest rate the metal sustains at $T = 300 \text{ K}$ without triggering twinning. Given that multi-junctions are markedly stronger than the regular binary junctions and can contribute a great deal to strain (forest) hardening, an interesting topic for further study is why the fraction of multi-junctions increases with increasing rate.

We relate the upward deviation in MD data for the saturated flow stress from a straight line of slope $1/2$ (Fig. 5b) to the presence of multi-junctions increasing with the strain rate. Normally taken to be constant, the Taylor hardening coefficient α can be corrected upward from its low-rate experimental value of $\alpha = 0.3$ (ref. 7) by a factor that accounts for the fraction of strong multi-junctions observed at each strain rate in our MD simulations (inset in Extended Data Fig. 7). Plotted in Fig. 5b as a red line is the so-corrected Taylor equation, based on a simple geometric mixing rule summing together contributions of binary and ternary junctions to the flow stress. The relatively narrow ranges of the flow stress and flow density explored in our simulations and our use of logarithmic coordinates make it difficult to see that the two lines in Fig. 5b are not parallel. Yet, when extrapolated to much lower quasi-static strain rates (dislocation density of about 10^{12} m^{-2}), at which the fraction of multi-junctions is expected to be nearly zero, the correction factor becomes unity and the two lines should intersect.

Perhaps the simplest and certainly most common characteristic of dislocation line geometry is the so-called character angle, or simply the character, defined as the angle between the Burgers vector of the dislocation and the direction of its line tangent. Disregarding the sign of its cosine, the character angle varies between 0° (pure screw character) and 90° (pure edge character). Dislocation characters between these two extremes are often referred to as mixed. In Extended Data Fig. 8a we show several histograms of the cosine of the character angle computed for a series of dislocation configurations encountered in the same straining simulation at rate $\times 25$. Extended Data Fig. 8b depicts the same evolution in terms of the ratio of near-screw ($0.99 < \cos(\theta) < 1.00$) to near-edge ($0.00 < \cos(\theta) < 0.01$) characters, with letters marking the instances where the character angle distributions shown in Extended Data Fig. 8a were computed. Because the character distributions are sharply peaked, the histograms are plotted on a base-10 logarithmic scale along the vertical axis. The histograms are normalized so that, when integrated over the cosine of the character angle from $\cos(\theta) = 0$ (pure edge) to $\cos(\theta) = 1$ (pure screw), the integral of the exponentiated histogram is equal to the total length of all dislocations contained in the simulation volume. Being equal to the product of dislocation density and simulation volume, the total length of dislocations increases rapidly through the transient and then saturates on reaching strain about 0.4.

The distribution is sharply peaked at zero cosine (near edge character) in the initial configuration A because our initial sources are hexagon-shaped prismatic loops with the normal vectors of the loop plane aligned along the Burgers vector of each loop. Therefore, in configuration A, the ratio of near-screw to near-edge characters is zero. On approaching the upper yield point, the loops extend into long screw dipoles, as reflected in a marked change in the character distribution where

the edge peak vanishes altogether but another sharp peak appears near the screw character. For instance, in configuration B, the ratio of near-screw to near-edge characters reaches 397. By comparison, in a random (uniform over the unit sphere) distribution of line orientations, shown as the dashed line on the right frame, the same ratio is equal to unity. This rapid redistribution from 'all edges' to 'mostly near-screw' results from the anisotropy of dislocation mobility that is common to body-centred-cubic metals; mobility of the screw dislocations is much lower than that of any other character, provided the temperature and/or dislocation velocities are not too high. At and past yield, the dislocations begin to collide, react and form a junction network in which most lines are seen to bend (bow) between two ends tied at junctions, manifested in a notable spreading of the screw peak. Although decreasing from its value at yield, the ratio of near-screw to near-edge characters remains high (at 86 in configuration C). At and beyond saturation strain about 0.4, the distribution of dislocation characters itself becomes stationary.

Although junction networks produced in our high-rate straining simulations are very dense and the flow stress is very high, forcing the dislocations to bow out considerably, we observe the screw character to remain dominant over the entire range of strain rates that we explored.

Effects of model size on simulation results. MD simulations, just as numerical simulations of any kind, are constrained by currently available computational resources, making it necessary to ensure that the dimensions of the simulation volume are at least minimally sufficient for the simulation results to be representative of the property or behaviour of interest (bulk metal plasticity in our case), even if simulations of the kind we report here are already extremely expensive and require massively parallel computing on unprecedented scales.

In Extended Data Fig. 10 we show the stress-strain and density-strain curves obtained in simulations performed under strain rate $\times 1$ ($1.11 \times 10^7 \text{ s}^{-1}$; midrange of the strain rates we explored). The thick solid black lines in both figures represent an MD simulation performed in the 1:2:4 brick-shaped periodic box with about 268 million atoms, precisely eight times greater than our standard simulation volume of about 33 million atoms. Data from the standard-sized simulations are also presented in the same figure for comparison. To evaluate statistical variations in the predicted strain response, we also performed eight independent standard-sized simulations, each of which was initialized to contain 24 prismatic dislocation loops (sources), but that were differently positioned in the simulation volume. Our simulation with 268 million atoms predicts a strain response that is identical, within statistical variations, to that of the standard-sized model. Furthermore, it appears that the simulated plasticity response is self-averaging in the sense that the stress-strain and density-strain curves obtained by averaging over eight smaller simulations are essentially the same as the corresponding curves obtained in the simulation of eight times greater size.

Deciding whether the simulation volume is sufficiently large for dynamic simulations of metal strength is a delicate issue. Simulation volumes that appear to be sufficiently large under one set of strain conditions may be inadequate under other conditions. Straining rate dependence is a relevant example. Many of our initial straining simulations were performed at rate $2.7 \times 10^8 \text{ s}^{-1}$ using a model of tantalum with 33 million atoms, which appears to be adequate—not too noisy, a good number of dislocations, and so on—for straining simulations at this specific rate (and higher). Yet the same volume appears to be a bit small for straining at rate $1.12 \times 10^7 \text{ s}^{-1}$, owing to a considerably lower dislocation density reached at this lower rate (as seen from Extended Data Fig. 10, the simulation size does not seem to affect the averages, but the simulated response in the smaller volume is rather noisy). To reduce statistical variations, it was desirable to increase the simulation volume for simulations at this rate; this is why we eventually ran the $1.12 \times 10^7 \text{ s}^{-1}$ simulation in the larger, 268-million-atom box.

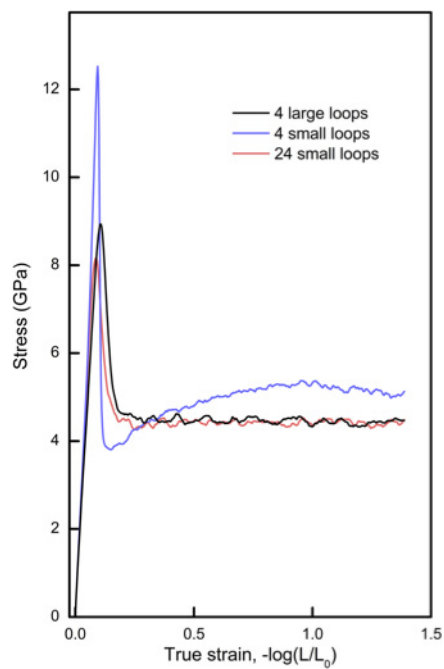
On the basis of experience, our recipe for selecting the system size boils down to the following considerations: we would like each dislocation to experience on average multiple (about 20) collisions and to not travel more than one box length over its lifetime in a simulation. The frequency of dislocation collisions and the distance each dislocation travels in a single simulation both depend on the resulting dislocation density. Because the latter is not known in advance, experimentation and adjustments to the simulation volume size may be required. For an MD simulation to conform to the just-mentioned requirements, the number of dislocations should be no less than 5,000. Given that dislocation density decreases with decreasing rate, any reduction in the strain rate, in addition to having to integrate the MD equations of motion over longer time intervals, generally requires larger simulation volumes, which increases the computational cost still further. For example, at the much lower dislocation densities of about 10^{12} m^{-2} that are typical of quasi-static straining experiments, a straightforward estimate for a cube-shaped simulation volume that is sufficiently large to contain 5,000 dislocations yields a size of about $17 \mu\text{m}$ (around 10^{15} tantalum atoms). Given that it would take 10^{15} time steps to integrate the MD trajectory to 1 s of simulated time, direct MD

simulations of such magnitude would require computational resource, measured for example in the units of atoms-steps, about 11 orders of magnitude greater than we have expended on our largest simulation performed at rate 10^6 s^{-1} .

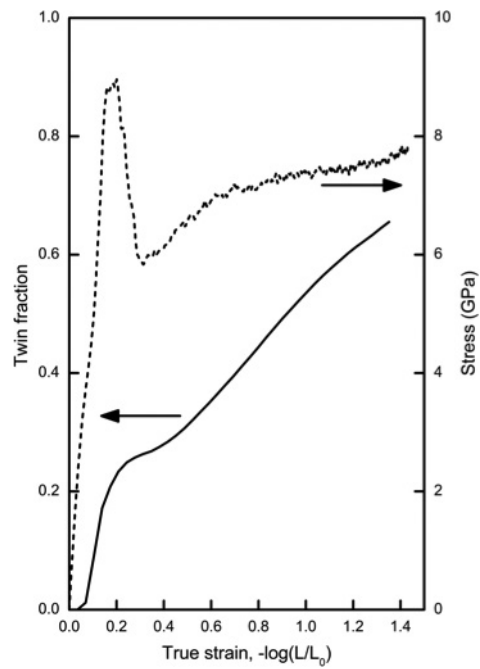
We continue to be mindful that size effects are possible at still larger scales. This is because elastic interactions are fundamentally scale-free (polynomial) and there may not be any characteristic length scale present in the system. Therefore, dislocation microstructure in a single crystal may well prove to be scale-invariant or fractal³⁸. In the context of numerical simulations of dislocation ensembles, no size may be large enough to make simulation results completely size-independent in a strict sense.

Data availability. The data that support the findings of this study are available from the corresponding author on reasonable request.

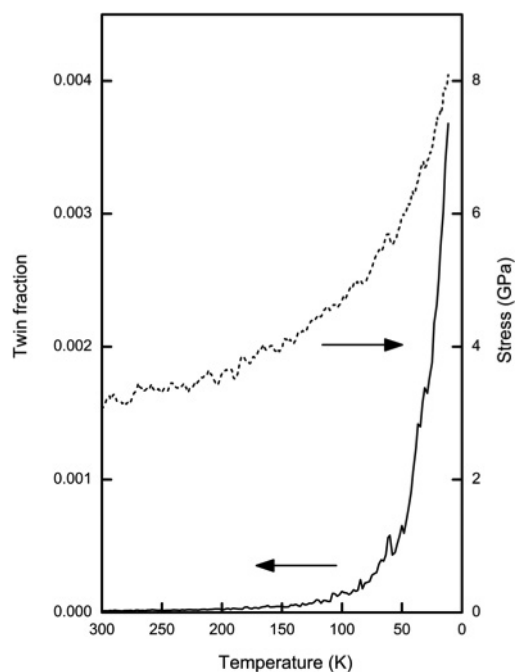
33. Plimpton, S. Fast parallel algorithms for short-range molecular dynamics. *J. Comput. Phys.* **117**, 1–19 (1995).
34. Li, Y. H. *et al.* Embedded-atom-method tantalum potential developed by force-matching method. *Phys. Rev. B* **67**, 125101 (2003).
35. Boxall, A. IBM's Sequoia tops the world's fastest supercomputer list. *Digital Trends* <http://www.digitaltrends.com/computing/ibms-sequoia-tops-the-worlds-fastest-supercomputer-list/> (2012).
36. Nguyen, L. D. & Warner, D. H. Improbability of void growth in aluminum via dislocation nucleation under typical laboratory conditions. *Phys. Rev. Lett.* **108**, 035501 (2012).
37. Mauro, J. C. & Smedskjaer, M. M. Unified physics of stretched exponential relaxation and Weibull fracture statistics. *Physica A* **391**, 6121–6127 (2012).
38. Hähner, P. *et al.* Fractal dislocation patterning during plastic deformation. *Phys. Rev. Lett.* **81**, 2470–2473 (1998).



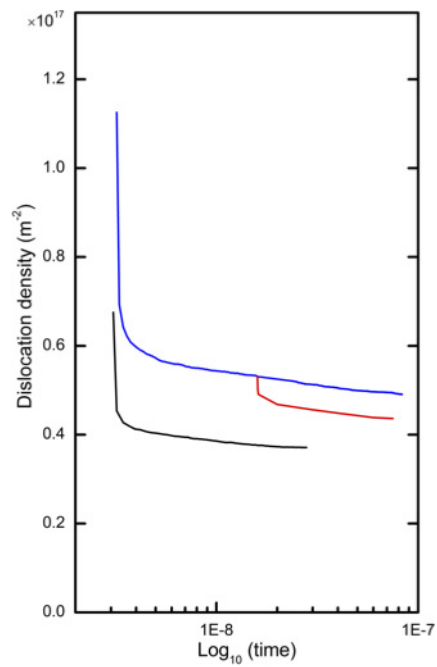
Extended Data Figure 1 | Plastic yield response depends on the initial density of dislocation sources. Stress as a function of true strain and specimen size computed in three MD simulations of compression at rate $\times 25$ from three different initial configurations of dislocation sources.



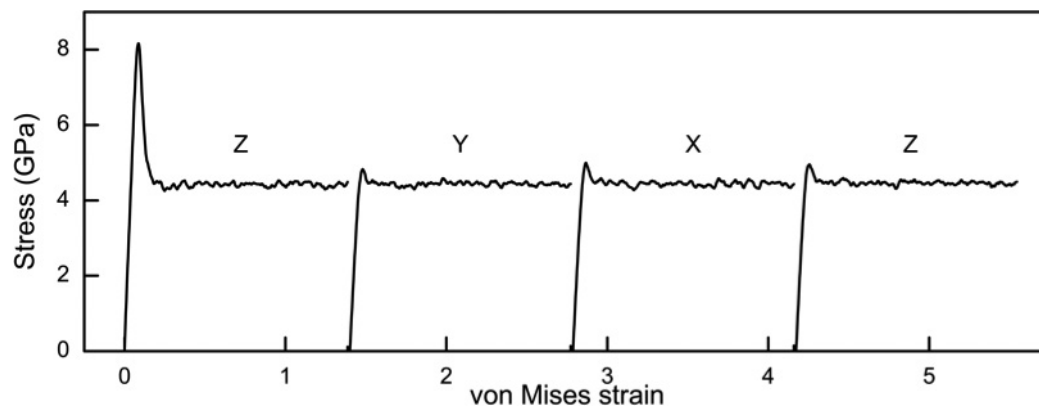
Extended Data Figure 2 | Detecting twinning during straining simulations. Stress (dashed curve) and the volume fraction of twins (solid curve) as a function of true strain under compression at rate $\times 5$ and temperature 25 K. The twin fraction was computed using the GSA.



Extended Data Figure 3 | Continuous cooling during straining permits detection of a twinning transition. Volume fraction of twins (solid curve) and flow stress (dashed curve) as a function of temperature, computed in a simulation at fixed rate $\times 5$ in which temperature was reduced at a constant rate from 300 K to 10 K. A twinning transition is identified by the temperature at which the twin fraction begins to rise rapidly from zero. Preceding this simulation, the crystal was pre-strained at the same rate $\times 5$ and a fixed temperature $T = 300$ K, where it attained a steady flow stress of 3.2 GPa.

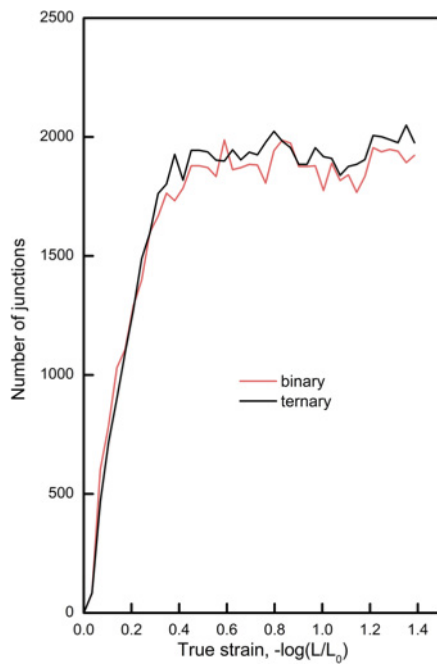


Extended Data Figure 4 | Relaxation of dislocation density after unloading of two crystals pre-strained at rate $\times 75$. The blue line represents isothermal relaxation after isothermal straining at $T = 300$ K. The red line depicts additional relaxation after two opposite surfaces of the simulated crystal were exposed to vacuum. The black line is adiabatic relaxation after adiabatic straining. Both isothermal and adiabatic relaxation simulations start at 2.5 ns in the end of isothermal and adiabatic pre-straining simulations (not shown).

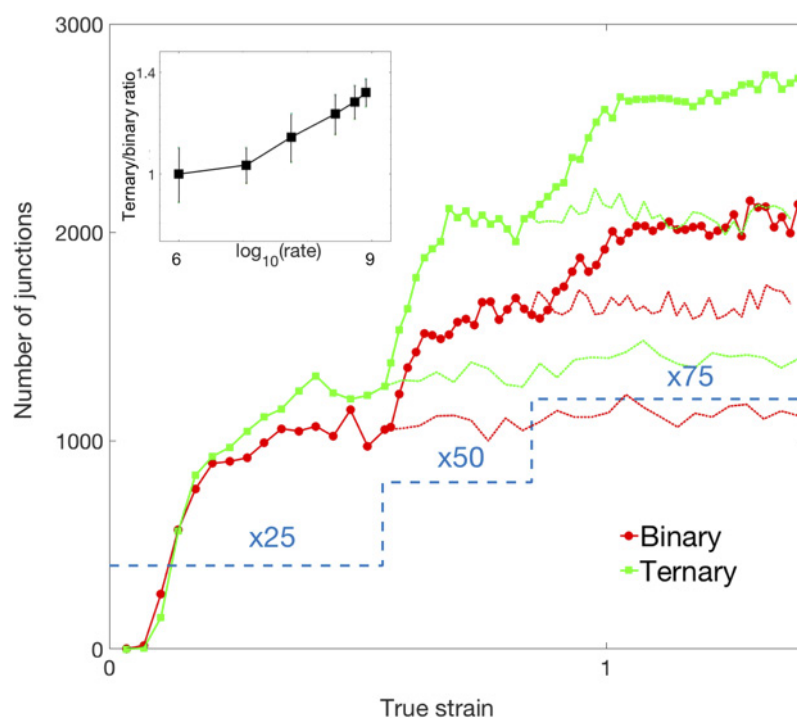


Extended Data Figure 5 | Kneading the metal. Stress as a function of true (von Mises) strain, computed under compression at constant true rate $\times 25$ along the three principal axes of the crystal. After compressing the

crystal to one-quarter of its initial length along the z axis, the strain axis is changed from z to y , from y to x , and then from x back to z . Letters above the stress-strain curves label the axes for each compression cycle.

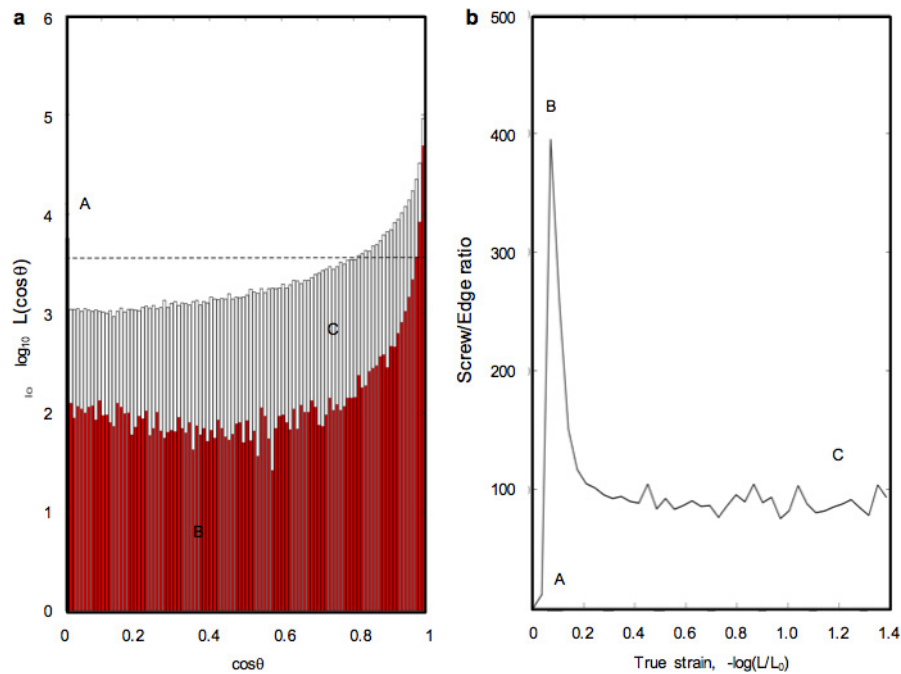


Extended Data Figure 6 | Evolution of dislocation network topology under compressive strain at rate $\times 1$. Following rapid dislocation multiplication at yield, regular binary junctions appear first (red line) closely followed by ternary multi-junctions (black line). After reaching stationary flow at a strain of about 0.4, dislocation density and network composition (topology) remain stationary within statistical noise.



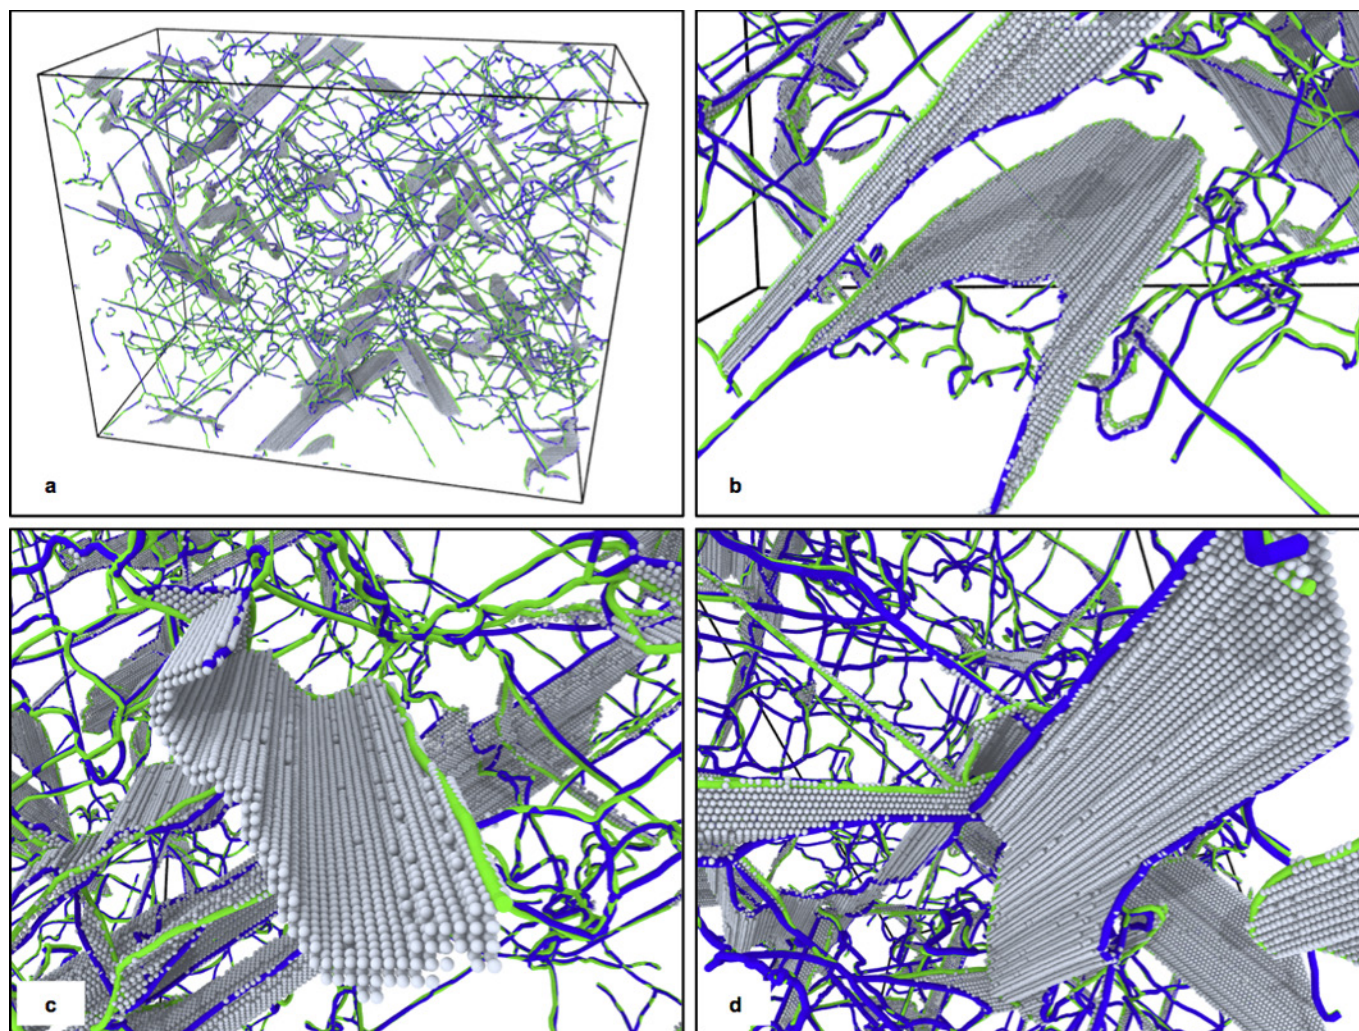
Extended Data Figure 7 | Network evolution along a stepwise increase in the strain rates. The number of binary junctions along a stepwise strain trajectory is shown as red circles and the solid red line, and the number of ternary junctions is shown as green squares and the solid green line. The dashed blue line shows the stepwise strain trajectory with strain rates marked at each rate step. The thin solid lines show the numbers of binary

(red) and ternary (green) junctions along continuations of the interrupted straining steps. The inset shows the ratio of the number of ternary junctions to the number of binary junctions attained in the saturated flow state as a function of strain rate. Error bars are the standard deviation from the mean values.



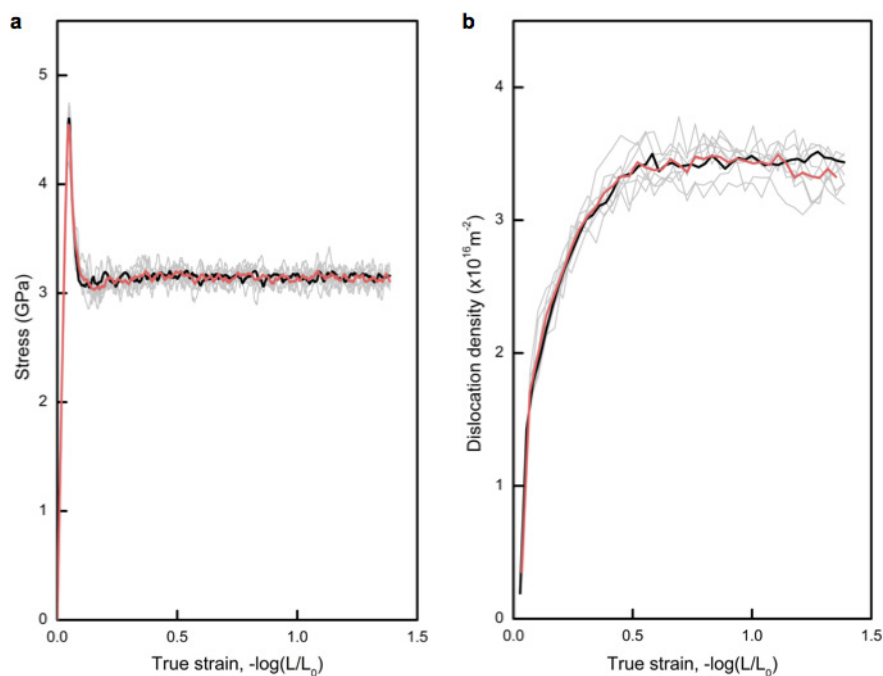
Extended Data Figure 8 | Evolution of dislocation characters under compression at rate $\times 25$. **a**, Histograms of dislocation character angle distributions computed for configurations A (black bar near zero), B (red bars) and C (open bars), which are marked along the strain trajectory in **b**. For reference, the dashed line depicts a hypothetical uniform distribution

of character angles with the same integral length of dislocations as in configuration C. The histogram counts are over the bins along the $\log[\cos(\theta)]$ axis. **b**, Ratio of the total length of near-screw dislocations to the total length of near-edge dislocations as a function of strain.



Extended Data Figure 9 | Differential slip trace analysis reveals how dislocations move and interact. **a**, The blue lines show positions and shapes of the dislocation lines in the initial configuration and the green lines show dislocations in the final configuration attained a few picoseconds later. The grey 'slip traces' consist of atoms whose local von Mises shear strain accumulated between the initial and the final dislocation positions exceeds the 0.15 threshold. Only a relatively small fraction of dislocations had swept substantial areas, whereas positions of most other dislocations in the two configurations coincide, suggesting little or no motion over the time interval. **b**, A magnified fragment of the same differential plot showing grey areas swept by several dislocations

in more detail. The smooth step seen on the slip trace in the foreground reveals a 'jog' (a turn of a dislocation line inside a crystal) on the moving dislocation. **c**, Cross-slip of a screw dislocation from its initial position (blue) to its final position (green). The shape of the cylindrical traced area reveals the detailed trajectory of the screw dislocation between its initial and final positions. **d**, Annihilation of two dislocations, as evidenced by a slip trace area bounded on its two sides by two blue lines: one straight screw dislocation above and one curved dislocation below. That annihilation has taken place is deduced from the absence of green lines, which would otherwise show final positions of the two dislocations.



Extended Data Figure 10 | Effect of simulation volume size on the strain response. **a**, Stress–strain response under straining at rate $\times 1$ ($1.11 \times 10^7 \text{ s}^{-1}$). The thick red line is the stress–strain response simulated in a volume eight times greater (about 268 million atoms) than the one used in most other simulations. The thin grey lines correspond to eight

independent simulations at the same rate, but performed in the standard-sized volume with about 33 million atoms. The thick black line was obtained by averaging over these eight simulations. **b**, The corresponding density–strain curves, with line colours and types matching the stress–strain curves on the left.