Tutorial Session 8d: T72S01- Basic Plasticity: Part 4
Relates to T72S01 Knowledge & Skills 1.13
Last Update 24/3/14
Incremental plasticity; Convexity of the yield surface; Normality of the plastic strain
increment; Flow rules; Reverse yielding: kinematic hardening versus isotropic
hardening and reality; Hysteresis loops, shakedown, ratcheting, residual stresses

Qu.:   Is the stress-strain curve the same in tension and compression?

Providing that true stress is plotted against true strain, the tension and compression curves for virgin material can often be very similar (though not invariably). If engineering stress were used, the reduction of the cross section in tension, and the increase of the section in compression, would lead to different curves.

Qu.:   What is the Bauschinger effect?

If a sample of material has yielded in tension, the compressive stress subsequently required to cause reversed yielding is smaller in magnitude than it would have been in the virgin state. In other words, the tensile yield has increased (hardening) but the compressive yield has decreased. This is the Bauschinger effect. Most (all?) polycrystalline materials are Bauschinger.

Qu.:   What is isotropic hardening?

Isotropic hardening is the name given to the following behaviour. If a material has been work hardened in tension, so that the yield strength in tension has increased from its virgin value σ_0 to some hardened value $\sigma_h > \sigma_0$, then unloading followed by re-loading in any stress state (compression, shear, whatever) will cause yielding only when the equivalent stress reaches σ_h. By "equivalent stress" we mean either Mises or Tresca, whichever criterion is applicable. Thus the yield strength in compression will be at a stress of $-\sigma_h$, and yield in shear will be at a shear stress of $\sigma_h/2$ for Tresca or $\sigma_h/\sqrt{3}$ for Mises. In deviatoric stress space, the yield surface has expanded uniformly in all directions (isotropically) by a factor of σ_h/σ_0. Thus, the Mises circle in the deviatoric plane has increased in radius from $2\sigma_0/3$ to $2\sigma_h/3$.

A Bauschinger material therefore does not harden isotropically. Hence most (all?) of the structural steels we deal with, i.e., polycrystalline materials, do **not** isotropically harden.

Qu.:   What is kinematic hardening?

Kinematic hardening is an extreme example of Bauschinger behaviour. Suppose a material has been work hardened in tension, so that the yield strength in tension has increased from its virgin value of σ_0 to some hardened value $\sigma_h > \sigma_0$. In kinematic hardening the yield surface in the deviatoric plane remains the same size and shape but is displaced to accommodate the hardening. This means that the yield strength in compression will now be *less* (in magnitude) than for the virgin material. Specifically, the yield stress in compression will be at a stress of $\sigma_h - 2\sigma_0$. (This is impossible if $\sigma_h > 2\sigma_0$, which exposes a limit on possible kinematic behaviour).

The shear yield will differ according to the orientation of the shear, with some orientations showing an elevated shear yield strength and some orientations a reduced shear yield strength.

Note that a kinematically hardened material is no longer isotropic after first yield.

Qu.: How do real materials harden?

Structural steels harden in a manner which is intermediate between isotropic and kinematic hardening. See diagram at end of these notes. There is no universal formulation. For austenitic steels some workers in BE have used a Lemaitre-Chaboche mixed hardening model.

Qu.: When do isotropic and kinematic hardening give different results?

If the loading is proportional (see below) then all hardening behaviours are the same since the directions of the stress and strain increments never change. Hence, different hardening behaviours can produce different predictions only for non-proportional loading.

Qu.: What is proportional loading?

Loading is proportional if the relative magnitudes of the stress components remain unchanged during loading, and the same is true of the plastic strain components after yielding. In other words, proportional loading involves moving monotonically along a radial line in deviatoric stress space (and the normal to the yield surface at this point remaining in a fixed direction as hardening occurs, i.e., as the surface moves outward).

Qu.: When is loading not proportional?

Examples of non-proportional loadings are: (a)if loads are reversed; or, (b)if one type of load is followed by another type of load. In the latter case, the two loads will generally produce stresses in different directions in deviatoric stress space, and hence the loading becomes non-proportional when the second load is applied.

Qu.: Is isotropic or kinematic hardening more conservative?

It depends upon the type of loading and the failure or degradation mechanism. Kinematic hardening materials will be less tolerant of cycling stresses, being more likely to fatigue and more prone to ratcheting (because larger plastic strains will be developed on reverse yielding and on re-loading). However, for a loading consisting of applied strains, e.g., welding residual stresses, isotropic hardening will result in larger stresses. Hence, isotropically hardening materials would be more prone to reheat cracking, other things being equal.

Qu.: What is a 'flow rule'?

A flow rule defines the magnitude and direction of the plastic strain increment for a given (infinitesimal) increment of all the stress components when yielding is taking place.

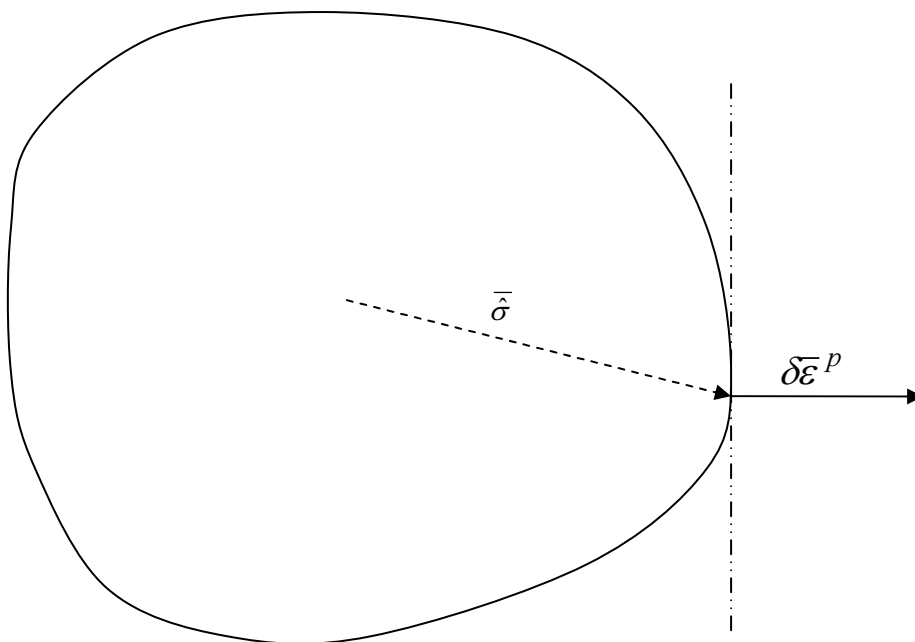**Qu.: What is meant by "the direction of plastic flow"?**

The plastic flow direction is the direction of the increment of plastic strain in deviatoric space, $\delta\varepsilon_{ij}^{p}$ - in other words the relative magnitudes of the components of $\delta\varepsilon_{ij}^{p}$.

Qu.:   How is the plastic flow direction determined?

The direction of plastic flow is determined by the Normality Rule. This states that $\delta\varepsilon_{ij}^{p}$ is normal to the yield surface at the point in question. Thus, for example, for a Mises material which hardens isotropically, $\delta\varepsilon_{ij}^{p}$ is radial at the current point on the Mises circle. For a kinematically hardened material, the normal to the displaced Mises circle also gives the direction of $\delta\varepsilon_{ij}^{p}$, but this will no longer be radial from the origin. For a Tresca material, $\delta\varepsilon_{ij}^{p}$ is normal to whichever side of the hexagon the stress 'vector' currently lies upon.

Qu.:   Where does the Normality Rule come from?

A general proof is provided in the Appendix. Here we derive normality from a plausible, if not rigorous, maximum work principle. Suppose a certain plastic strain increment, $\delta\varepsilon_{ij}^{p}$, is imposed on a body by an external agency. If we imagine that the material resists the deformation as much as it is able, then the stresses which result will maximise the work which the external agency has to do to impose the deformation, $\bar{\hat{\sigma}}\cdot\delta\bar{\varepsilon}^{p}$. (We have used a vector notation here since vectors in the principal deviatoric plane will suffice). But $\bar{\hat{\sigma}}$ lies on the yield surface, and hence $\bar{\hat{\sigma}}\cdot\delta\bar{\varepsilon}^{p}$ is maximised by choosing $\bar{\hat{\sigma}}$ so that $\delta\bar{\varepsilon}^{p}$ is normal to the surface, since this maximises the component of $\bar{\hat{\sigma}}$ in the direction of a given $\delta\bar{\varepsilon}^{p}$, thus,

Any other placing of a given $\delta\bar{\varepsilon}^{p}$ on the yield surface would lead to a small value of the work $\bar{\hat{\sigma}}\cdot\delta\bar{\varepsilon}^{p}$. Hence,
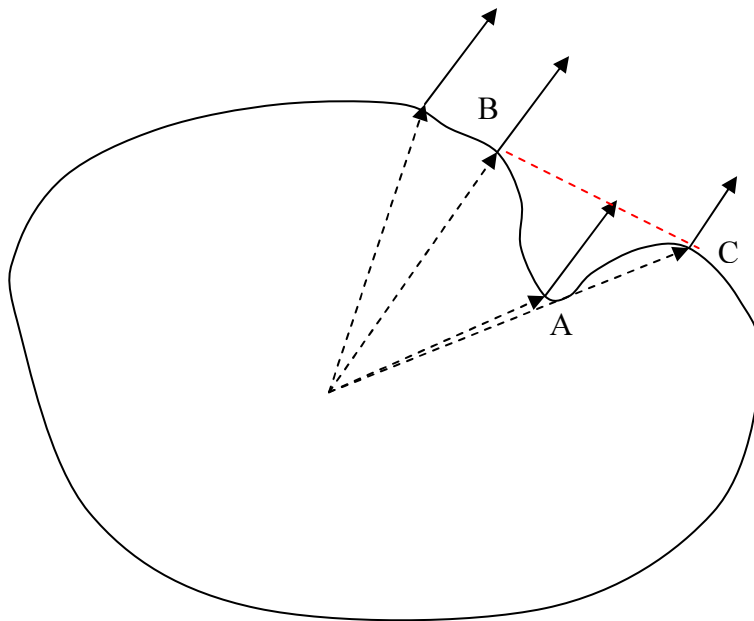
maximum plastic work implies normality

No.

The Convexity Rule says it must be convex. (A convex closed curve is such that the tangent at any point lies entirely outside the curve, except at the point of tangency). Again there is a general proof of convexity in the Appendix.

Here we note only that convexity is necessary given the normality rule if any given $\delta \overline{\varepsilon}^{\,p}$ is to correspond to a unique value of stress, $\overline{\hat{\sigma}}$. A yield surface with a re-entrant part would give rise to an ambiguity of stress state for a specified flow, as illustrated by the diagram which follows:-

A non-convex yield surface would have multiple possible stress states, for a given plastic flow direction, $\delta \overline{\varepsilon}^{\,p}$, all obeying the normality rule.

Actually, even a convex surface, part of whose boundary is flat, will have an ambiguity in stress state for a given strain increment. The Tresca hexagon is an example. However, all points on a flat correspond to the same amount of work, $\overline{\hat{\sigma}} \cdot \delta \overline{\varepsilon}^{\,p}$. In contrast, the re-entrant point A, above, corresponds to less work, $\overline{\hat{\sigma}} \cdot \delta \overline{\varepsilon}^{\,p}$, than the points B and C. If we are guided by the principle that the material will resist to the maximum, then point A will be unstable and the material will flip preferentially to B or C. The same argument applies for all points within the re-entrant 'dimple'. So the material will never rest within the re-entrant part. In effect we can flatten off the dimple, as indicated by the red dashed line. In this way the yield surface necessarily becomes convex. In this sense we can also say that,

maximum plastic work implies convexity

Since the Appendix claims to prove normality and convexity without the use of the maximum plastic work principle, the above argument can be turned around to be a demonstration of the maximum work principle (i.e., that, for a given stress state, plastic flow will be in a direction so as to maximise plastic work).

Qu.:   What is a hysteresis loop?

A hysteresis loop is the locus of the stress-strain points which is generated by a sequence of loads which start and finish at the same loading condition. Repeating the same sequence of loading conditions results in a set of hysteresis loops which may or may not over-plot. Such loops could be plotted for any components of stress and strain, but equivalent stress and strain are most often employed. However, there is an important subtlety in this case because equivalent quantities cannot be negative.

To overcome this, a "signed equivalent stress" is often used, defined as the Mises stress times the sign of the principle stress which has the greatest magnitude (which can be a bit messy, introducing discontinuities). Similarly, the plastic equivalent strain increment is really always a positive quantity, and hence the plastic strain always increases monotonically. This is what ABAQUS gives as PEEQ, for example. But it is convenient to define effective plastic strain increments which are signed, and which thus allow the strain to reduce as well as increase (without which there could be no closed hysteresis loop!). Again, the sign of the effective plastic strain increment can be defined by the sign of the largest principle plastic strain component. By this means, closed, or nearly closed, hysteresis loops can be generated by a finite element analysis from equivalent quantities. R5V2/3 also uses signed equivalent quantities for constructing stress-strain hysteresis loops.

Qu.:   What does the area of a hysteresis loop represent?

Assuming the hysteresis loop plots equivalent stress and strain, its area is the plastic energy dissipated per unit volume within the material during the loading cycle.

Qu.:   If load is applied and then removed, does the stress return to zero?

In general, no.

Clearly it does if the material remains elastic. Also, in the case of uniaxial tension, the net stress will return to zero. However, in general, the stress at an arbitrary point of a body which has yielded and then had the load removed will not be zero. Take for example, a structure with a local stress concentration such as a notch. Suppose the elastically calculated stresses would be several times the yield stress. Upon unloading the material near the notch will probably yield in compression. The structure will be left with residual stresses, in this case large and compressive local to the notch (with smaller balancing tensile stresses away from the notch).

Qu.:   What is shakedown?

If a loading-unloading sequence is repeated over and over, other than purely elastic behaviour three generic behaviours are possible: either, (a)the stress-strain trajectories at all points of the structure settle down to move up and down the elastic line; or, (b)the stress-strain hysteresis loops persist but settle down to repeat the same loop; or, (c)neither of these, i.e., the hysteresis loops persist and do not over-plot. In case (a) the structure is said to 'shakedown'. Thus, despite the rather misleading name, shakedown is a desirable state of affairs. It means that the structure is essentially elastic after some initial plasticity during the first few load cycles.

Ratcheting is case (c), above. When a repeated load-unload cycle does not produce a repeating hysteresis loop after a few cycles, it will usually be because the loop is progressing steadily to the right, i.e., accumulating more and more plastic deformation on every cycle. Ratcheting is a serious structural threat, indicative of excessive cyclic stresses. Ratcheting should not occur in a well designed structure. It is likely that a ratcheting structure will fail relatively quickly, if only because something outwith the design basis is occurring, though in principle the ratchet strains could be small and take a large number of cycles to fail. Ratcheting is the extreme opposite of shakedown.

Qu.:   Do all structures either shakedown or ratchet?

No.

Case (b), above, is the intermediate case: that of a stable hysteresis cycle ("stable plastic cycling"). This condition will cause fatigue damage to accumulate on each cycle, the amount being sensitive to the strain range of the hysteresis cycles. In principle even structures which shakedown may be subject to fatigue. But the fatigue endurance for strain ranges within the elastic regime (e.g., <0.2%) will generally be enormous. In contrast, when hysteresis cycles occur, the plastic strain range may be the order of several percent strain, which can cause rapid accumulation of fatigue damage.

The occurrence of hysteresis cycles, rather than elastic cycling, is even more significant when operation is in the creep temperature regime. In this case the hysteresis cycle can repeatedly re-set creep dwell stresses, leading to creep-fatigue damage accumulation on each cycle – a potentially virulent degradation mechanism and one which is often the most relevant to in-reactor components in the upper regions of AGR reactors and boilers.

In practice, it is those structures which lie between strict shakedown [behaviour (a)] and ratcheting [behaviour (c)] which are the chief subject of R5V2/3 crack initiation assessments. Confusingly, perhaps, R5 also introduces the concept of "global shakedown". This is a subset of behaviour (b) in which the regions exhibiting cyclic plasticity are confined to <20% of any section of the structure [the rest of the structure exhibiting behaviour (a)]. This is a precaution by R5 to ensure that its recommended procedures for assessing crack initiation are valid. These R5 procedures apply outside of strict shakedown but within global shakedown.

Qu.:   Can residual stresses be deduced from the deformation of the structure?

No.

A given plastic deformation can be achieved in a multitude of ways, in general corresponding to different stress states. For example, a bar may be bent to a given extent by loading and then unloading. However, the same bend may be achieved by initially applying too much load, then unloading and reverse loading enough to correct the over-bend – followed by unloading again. The latter procedure will result in residual stresses of opposite sign to the first case. This can happen only because plasticity is irreversible, and plasticity theory is incremental and exhibits load order dependence.
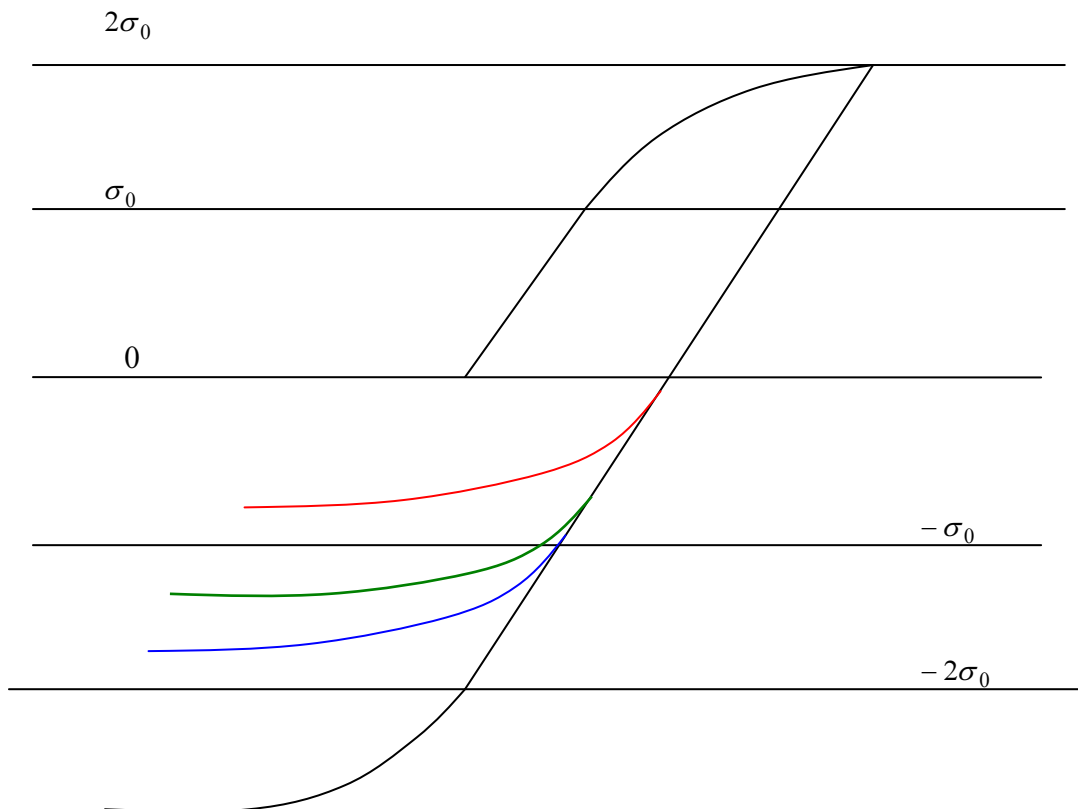
The direction of $\delta\overline{\varepsilon}^{\,p}$ is defined by the normal to the yield surface, which for the isotropic Mises rule is also the 'radial' direction, i.e., $\hat{\overline{\sigma}}$. The magnitude of a small plastic strain increment is proportional to the increment of Mises stress. Hence, so far, we have $\delta\overline{\varepsilon}_{ij}^{\,p} \propto \hat{\sigma}_{ij}\,d\overline{\sigma}$. The rest of the flow rule is defined by the uniaxial stress-strain curve, which can also be interpreted as an equivalent stress-strain curve. Thus, if we assume a power law hardening: $\overline{\varepsilon} = A\overline{\sigma}^{\,n}$, and hence $d\overline{\varepsilon} = An\overline{\sigma}^{\,n-1} \cdot d\overline{\sigma}$, then we get,

$$\delta\overline{\varepsilon}_{ij}^{\,p} = \frac{3}{2} \cdot An\overline{\sigma}^{\,n-2}\hat{\sigma}_{ij}d\overline{\sigma} \qquad (1)$$

The factor of 3/2 is to ensure the uniaxial behaviour is regained, since in this case $\hat{\sigma}_x = \frac{2}{3}\overline{\sigma}$. This is a simple explicit representation of a flow rule. But, for real hardening behaviour, when the yield surface moves, it is harder to specify the rule algebraically.

Schematic Example of Isotropic, Kinematic and Typical Hardening

Black reverse yielding curve = isotropic hardening
Red reverse yielding curve = kinematic hardening
Blue reverse yielding curve = intermediate hardening behaviour, but non-Bauschinger
Green reverse yielding curve = typical realistic hardening behaviour, Bauschinger

APPENDIX

Proof of Convexity & Normality in Plasticity

The cunning thing is to recognise that the agency causing the yielding may not be responsible for the whole state of stress.

Suppose a body is initially fully elastic but has various loads, residual stresses, etc acting, so that the state of stress at a certain point is σ_{ij}^0, within the yield surface. These stresses are in equilibrium with the applied loads, so that $\sigma_{ij,j}^0 + b_i^0 = 0$, where $\bar{b}^0$ is the force per unit volume applied by the initial loads.

Now an external agency, which we call the "yielding agency", causes additional stressing, $\Delta\sigma_{ij}$, and takes the stress to $\sigma_{ij} = \sigma_{ij}^0 + \Delta\sigma_{ij}$ which lies on the yield surface. The external agency now causes some plastic straining, $\delta\varepsilon_{ij}^p$. Because the yielding agency has applied only those forces which give rise to $\Delta\sigma_{ij}$, the work done by the yielding agency (per unit volume) is $\Delta\sigma_{ij}\delta\varepsilon_{ij}^p = \left(\sigma_{ij} - \sigma_{ij}^0\right)\delta\varepsilon_{ij}^p$.

The irreversibility of plasticity means that it is only possible for the yielding agency to do work on the material. It is not possible for any element of the material to do work on the external agency whilst undergoing plastic straining, since this would be synonymous with recovery. Consequently $\Delta\sigma_{ij}\delta\varepsilon_{ij}^p = \left(\sigma_{ij} - \sigma_{ij}^0\right)\delta\varepsilon_{ij}^p$ is everywhere positive indefinite, and this holds for an arbitrary choice of σ_{ij}^0 within the yield surface and σ_{ij} on the yield surface.

A little thought shows that this gives us both normality and convexity. The "vector" $\Delta\sigma_{ij}$ can be drawn from any point within the yield surface to end at any point on the yield surface, but it must always lie "behind" the strain increment "vector", $\delta\varepsilon_{ij}^p$, where "behind" means making an angle >90^o. This is only possible if both normality and convexity apply.

WARNING: I am not sure about the status of this 'proof'. I cannot see anything wrong with it, but I previously thought that normality & convexity could not be demonstrated with full mathematical rigor - requiring an irreducible degree of empirical input. Spot the flaw?