Chpater X How to build a Biological oscillator

Oscillations are thrilling – hearts beat, cells divide every cell cycle, circadian clocks keep time, neurons click in trains of regularly spaced spikes. Biological oscillations attracted theoretical work from pioneers such as Arthur Winfree, Albert Goldbeter and John Tyson, and are still an active area of research. Let's discuss the design principles of biological circuits that oscillate.

Oscillations require negative feedback and delay

At the heart of an oscillator is a negative feedback loop. Molecule x acts to reduce its own amounts, so that high levels go to low and then back to high and so on. Negative feedback on its own, however, is not enough. The simplest negative feedback motif - negative autoregulation $\frac{dx}{dt} = f(x) - \alpha x$, does not oscillate but instead monotonically returns to steady-state as we saw before (Fig 1). In order to oscillate, you need to add a sizable delay to the negative feedback loop.

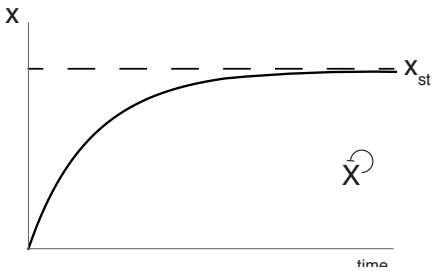

Figure 11.1

Negative feedback plus delay makes me vividly remember the shower we had when I was a child. The water started cold. I would turn on the hot water – acting as a feedback controller. But the hot water took some time to arrive, so I would turn the handle too much, and the water would be scathing- ouch! So I would turn it back strongly to the cold, but because of the delay I would go too far and the water would be freezing- Arrgh! And so one in a cycle of Ouch! Arrgh! Ouch! Fortunately, in our more modern shower, there is less delay and I can easily tune the desired temperature.

A delay in biological circuits can be achieved by adding components in the negative feedback loop to make longer paths in the circuit. Autoregulation is just a single self-closing arrow, a one-step path. Add another species, y, and you get a negative feedback loop made of two arrows (Fig 2). Here you start to see a hint of oscillations: you can get damped oscillations with pulses that settle down to steady-state (Fig 3). Damped oscillations require (i) strong (preferably cooperative) feedback and (ii) that the timescales on the two arrows are similar. If the timescales are very different, with one arrow much faster than the other, the fast path is not much of a delay element, and the circuit acts effectively like autoregulation with no damped oscillations (Fig 4).

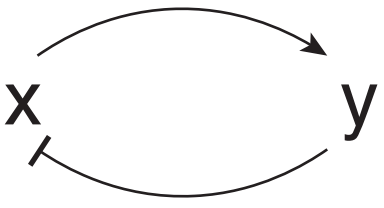

Figure 11.2

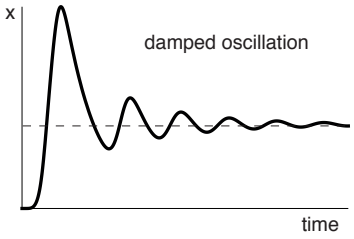

Figure 11.3

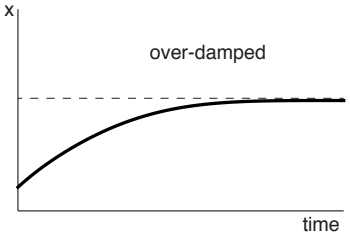

Figure 11.4

An easy way to grasp the roles of feedback strength and timescales is to use **linear stability analysis**, as in the solved exercise below. If you need to refresh your memory of linear stability theory, go to appendix LIN.

Solved Example X.1

Show that a 2-component negative feedback loop shows damped oscillations if timescales are similar enough and feedback is strong enough.

Consider a 2-node negative feedback circuit where x activates y according to the increasing function g(x) and y represses x according to the decreasing function f(y) (Fig 5)

$$\frac{dx}{dt} = f(y) - \alpha_1 x,$$

$$\frac{dy}{dt} = g(x) - \alpha_2 y$$

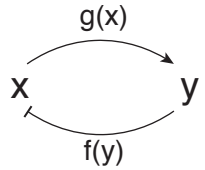

Figure 11.5

The α's are the removal rates that set the timescales of the two arrows. The dynamics of small perturbations x(t) and y(t) around the fixed point xst, yst are governed by the linear equations

$$\frac{dx}{dt} = \beta 1\, y - \alpha 1\, x \qquad \frac{dy}{dt} = \beta 2\, x - \alpha 2\, y$$

where the feedback parameters β_1 and β_2 are the derivatives of the g(x) and f(y) at the fixed point: $\beta_1 = \frac{df}{dy} < 0$, and $\beta_1 = \frac{dg}{dx} > 0$. Note that $\beta_1 = \frac{df}{dy}$ is negative because f is a decreasing function, whereas beta2 is positive because g is increasing. Thus, the dynamical system can be described in matrix notation by

$$\frac{d}{dt}\begin{pmatrix} x \\ y \end{pmatrix} = J \begin{pmatrix} x \\ y \end{pmatrix}, \qquad J = \begin{bmatrix} -\alpha_1 & \beta_1 \\ \beta_2 & -\alpha_2 \end{bmatrix}$$

where J is known as the **Jacobian matrix.** For a general system dx/dt=f(x,y), dy/dt=g(x,y), the jacobian is given by the derivatives evaluated at the fixed point,

$$J = \begin{bmatrix} \frac{\partial f}{\partial x} & \frac{\partial f}{\partial y} \\ \frac{\partial g}{\partial x} & \frac{\partial g}{\partial y} \end{bmatrix}.$$

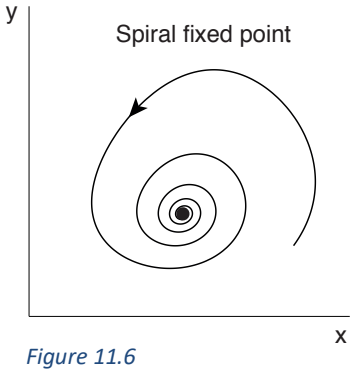

Figure 11.6

The dynamics of such linear equations are determined by the two eigenvalues λ_1 and λ_2 of the matrix J, because solutions can be written as a sum of exponentials of time: c1 exp(λ_1t)+c2 exp(λ_2t). Damped oscillations occur when the eigenvalues have an imaginary part $\lambda_{1,2} = a \pm ib$, due to Euler's formula $e^{(a+ib)t} = e^{at}(\cos(\mathrm{bt}) + \mathrm{i}\sin(\mathrm{bt}))$. This is an oscillating wave with frequency determined by the imaginary part $\omega = \frac{2\pi}{b}$. The amplitude of the wave decays exponentially to zero if a<0. The result is known as a **spiral fixed point**, because the system spirals down into a steady state solution (Fig 6).

To find the eigenvalues we solve the characteristic equation of the matrix J, $(-\alpha_1 - \lambda)(-\alpha_2 - \lambda) - \beta_1\beta_2 = 0$ to find $\lambda_{1,2} = (-(\alpha 1 + \alpha 2) \pm \sqrt{(\alpha 1 + \alpha 2)^2 - 4(\alpha 1\alpha 2 - \beta 1\beta 2)})/2$ The eigenvalues in this circuit always have a negative real part and so the steady-state is stable and all initial conditions flow back to it (a rule for stability of two-variable systems is if the sum of the

diagonals (the trace of the matrix) tau is negative and the determinant Delta is positive, which is the case here as can be seen from the sign structure of the Jacobian =[- +; + -]. Here tau=-(a1+a2)<0, and Delta= α1α2-β1β2>0). The eigenvalues have imaginary parts, and thus produce damped oscillations, when the term inside the square root is negative (stable spiral fixed point) namely

(α1-α2)^2<-4 β1β2

Note that the right hand side is positive since beta2<0. When the timescales of the two arms are equal (α1=α2), damped oscillations always occur, for any **feedback strength** |β1β2|. The frequency of the damped oscillations is given by ω=2 pi/ sqrt((α1-α2)^2-4 β1β2). The larger the mismatch in timescales |alpha1-alpha2|, the larger the feedback strength|β1β2| needed for damped oscillations. If feedback is not strong enough compared to the timescale separation, the system is **overdamped** and decays monotonically to the fixed point with no overshoot (Fig 4). Thus, strong separation of timescales counteracts the tendency to oscillate.

What exactly do I mean by feedback strength? The strength β1β2 is determined by the slopes of the regulation functions g and f at the steady state point, dg/dx and df/dy. The more steep these regulation functions- for example the higher their Hill coefficient - the stronger the feedback. Cooperativity enhances the tendency to oscillate. Notice that feedback strength is the product of the two betas, so that if one arm has beta=0, the circuit is effectively not a feedback loop at all, and total feedback strength is zero.

Cooperativity helped the oscillations in my childhood shower, because the faucet had a very steep curve- very hot or very cold for most of the range, making it harder to tune in on the right temperature.

==========

Thus, a two-step negative feedback loop can only show damped oscillations. This observation prompted Galit Lahav, when she was a postdoc in my group, to try to see these damped oscillations in living cells, using an important feedback loop that involves a protein known as the 'guardian of the genome', p53.

p53 is called guardian of the genome because it governs cell decisions when DNA is damaged. The cell decides to either repair the DNA, or, if it is too damaged, to avoid becoming cancerous by committing programmed cell death or becoming a zombie-like senescent cell that stops dividing. That is why p53 is mutated in most cancers, bypassing cell death and allowing cancer cells to proliferate despite damage. p53 forms a negative feedback loop with another protein mdm2: p53 transcriptionally activates mdm2, and mamd2 leads to the degradation of p53 (Fig 7).

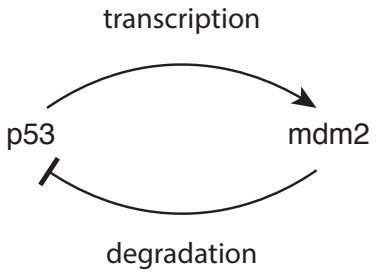

Figure 11.7

Galit Lahav fused the genes for p53 and mdm2 to green and red fluorescent proteins. That way, she could see in the microscope how red and green fluorescence varied in individual human cells over time, reporting for the changes in the two proteins. This was an advance over the way experiments on p53 had been done before, by averaging over millions of cells, and thus potentially masking out dynamic processes.

Galit gave the cells some gamma irradiation to induce DNA damage, and filmed the cells. She even brought a bed to the lab because she had to focus the microscope every 15 minutes over 24h of filming (this is heroic. A year later we got a microscope with automated focus).

To our surprise, we did not see damped oscillations, but instead full-fledged oscillations that do not damp out. p53 enters and exits the nucleus with pulses that have noisy amplitudes and precise 6h period (Fig 8). Mdm2 also oscillates, with the opposite phase.

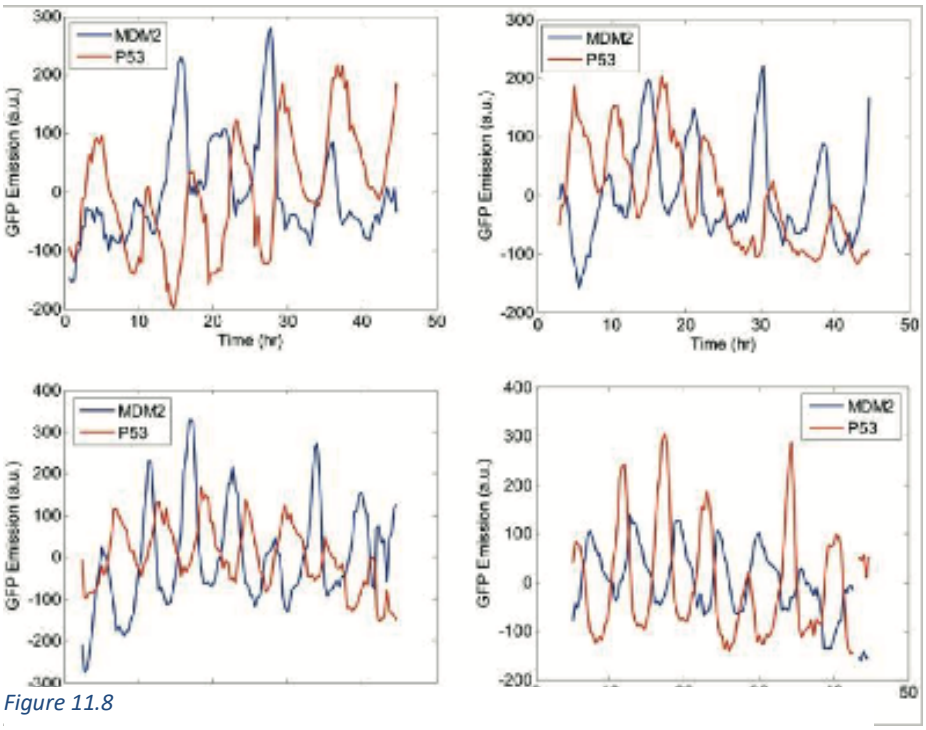

Figure 11.8

Noise can induce oscillations in systems that have only damped oscillations on paper

It took us a while to figure out what is going on. It turns out that even circuits that show only damped oscillations on paper, such as a two-component negative feedback loop, can still oscillate indefinitely in the cell. This occurs when noise is strong enough. Noise kicks the system away from the spiral fixed point and prevents the oscillations from damping out (Fig 9). It's like a damped spring that is constantly perturbed. As a result, the circuit shows pulses with noisy amplitudes but rather precise frequency. The frequency is that of the original (noiseless) damped oscillation, given by the imaginary part of the eigenvalues. Another way to think about this is that the damped oscillator has a resonance frequency, and therefore amplifies the part of the noise which has that frequency.

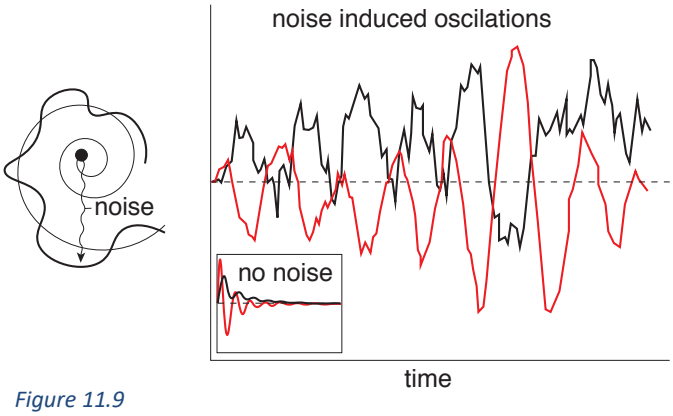

Figure 11.9

A diagnostic for such **noise-induced oscillations** is that the amplitudes of the pulses are more variable than their frequency (Fig 8), and that the amplitude increases with noise strength. Theoretical work (Lang PMC Biophy 2009) shows that the distribution of the pulse peak amplitude A goes as P(A)~A exp(-A^2/Ao^2) where Ao is the ratio of the noise amplitude and real part of the spiral-fixed-point eigenvalue[ref]. This formula describes the p53 pulses well (Dekl Alon 2010).

In recent years, several other TFs have been found to oscillate in and out of the nucleus. Some TFs, such as NFKb, show oscillation pulses with noisy amplitude and accurate frequency like p53. Sometime transcription factors have several isoforms, with one showing oscillations and the other a graded response to a given signal (such as the TFs NFAt1 and NFAt4 in immune signaling). Other TFs, such as Crz1, show trains of pulses of nuclear entry, whose frequency increases with the input signal, while their amplitude does not depend on signal. One reason that TFs may oscillate is to keep exciting downstream genes that would otherwise show exact adaptation to TF level, as described in the chapter on fold-change detection. Pulses 'wake up' circuits that otherwise adapt. Additional possible utility of TF oscillations is discussed in exercise XX.

Delay oscillators

Full-fledged, undamped oscillations even without noise can appear if we go to feedback loops with three or more steps. In order to oscillate, such loops need to have strong feedback and similar timescales for the different steps.

Figure 11.10

A three component negative feedback loop featured in one of the first theoretical models of biological oscillators, by Goodwin[]. Several decades later, a three component loop helped to inspire the field of synthetic biology, when Michael Elowitz and Stanislas Leibler in 2000 built a cycle of three repressors, called the **repressilator** (Fig 10). They linked three repressors in E coli, with a green fluorescent gene (GFP) as a readout. To make sure the circuit parameters supported oscillations, Elowitz and Leibler made the timescales of the components as similar as possible. The repressilator oscillated in E. coli, with the GFP readout blinking green, black, green with a period of about 8 hours. The repressilator was recently updated in a more minimal and precise version [Paulsonxx (2017)].

What happens when you add more than three components into the negative feedback loop? The more components in the cycle, the larger the range of parameters for oscillations, and the weaker the degree of cooperativity required for oscillations. Exact solutions are in exercises xx. The frequency of the oscillation in such **delay oscillators** is generally proportional to the sum of the half-lives of the components, the overall delay time.

Many biological oscillators have a coupled positive and negative feedback loop motif

To sum up so far, negative feedback plus delays and/or noise can provide oscillations. Nonlinearity (cooperativity) and similar timescales for the opposing arms help the feedback loop to oscillate. But when we look at the circuits for the best-studied oscillators in biology, such as heart cells, neurons and cell cycles, we see an additional feature - a *positive feedback loop* is added to the negative feedback loop (Fig 11). What is the role of positive feedback?

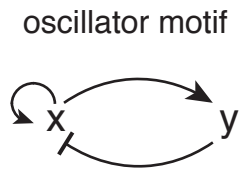

Figure 11.11

The positive feedback loop adds a delay as we saw in chapter 4, and delay helps oscillation. It increases the parameter range for oscillations. For example, positive feedback can make a two-node negative feedback loop show sustained oscillations even with one arm much faster than the other. Separation of timescales between the interactions in the negative feedback loop is, in fact, a recurring feature of the oscillator motif of Fig 11.

To see how positive feedback can make a two-component loop oscillate, we can use linear analysis of the fixed point. Without positive feedback, the two eigenvalues have negative real part and we have a stable fixed point or stable spiral. Positive feedback can make the real part go positive- turning the stable spiral into an unstable spiral (Fig 12).

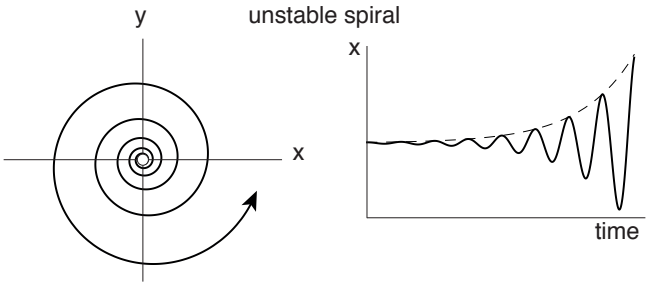

Figure 11.12

The trajectories spiral out – but can not diverge to infinity because, once concentrations rise sufficiently, all feedback terms saturate and we are left only with the removal terms –alpha x, -alpha y that push concentrations back down (biochemical circuits have the saving grace that concentrations cannot diverge and cannot go negative). Thus trajectories are kept somewhere away from the unstable fixed point and also away from infinity. A fundamental theorem of two-component dynamical systems (Poincare-Bendixon theorem, Storgatz(xx)) shows that such confined trajectories settle into a sustained oscillation (exercise XX) called a **limit cycle** (Fig 13).

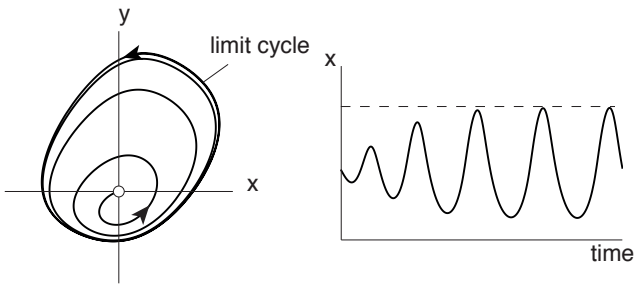

Figure 11.13

Example: positive feedback can destabilize the fixed-point of a two-component negative loop

We can modify the 2-node feedback loop analyzed above, by adding positive autoregulation to x. As a result, its production rate changes from f(y) to becomes a function of both x and y, f(x,y), which rises with x (autoregulation) but drops with y (negative feedback). The equations are

$dx/dt=f(x,y)-a1\ x$,

$dy/dt=g(x)-a2\ y$.

The Jacobian matrix at the fixed point in this case is $\begin{bmatrix} -\alpha 1 + P & -\beta 1 \\ \beta 2 & -\alpha 2 \end{bmatrix}$ where $P = \frac{\partial f}{\partial x} > 0$ is positive due to positive autoregulation. The negative feedback arms are -b1=df/dy<0, b2=dg/dx>0. The stability of the fixed point is determined by the real part of the eigenvalues. The sum of the real parts is equal to the trace of the matrix (the sum of the diagonal terms), tau=P-alpha1-alpha2. When tau becomes positive, one of the eigenvalues has a positive real part and the fixed point becomes unstable. This occurs when positive autoregulation strength exceeds removal, P>Pc=(alpha1+alpha2). The sign structure of the Jacobian goes from the stable [- +, - -] to [+ +,- -]. If the positive feedback P is increased slowly, a stable spiral fixed point turns unstable when P=Pc in what is known as a **Hopf bifurcation**.

==========

Positive feedback can also make a more dramatic contribution: bistability. Bistability is powerful in an oscillator, because it makes the oscillations more decisive and less noisy. The circuit makes sharp transitions between two states of the fast variable, going tic-toc between high and low concentration. The amplitude is well-defined by the difference between these states, and hence frequency can be changed if needed without affecting the amplitude. The role of bistability was worked out nicely in one of biology's most fundamental oscillators, that drives the cell cycle.

Cell-cycle circuits are usually complicated, with dozens of components that act as checkpoints to make sure important steps such as replicating the DNA are completed before cells divide. In some cells, however, the cell-cycle circuit is stripped down to a minimum, offering a model system for basic understanding. An example is the circuit in charge of the first divisions of the frog egg, which occur every 20 min (Fig 14). The circuit has a negative feedback loop between X and Y and positive feedback on x (X is CDK1, a kinase that activates many proteins for cell division, including Y, called APC which degrades active X). X can be in two states – phosphorylated xp, the active form and Xo, the unphosphorylated form of X, which is inactive. Xp acts as a kinase that can phosphorylate and thus activate proteins for cell division. Xp also activates Y which degrades X, forming the negative feedback loop. A protein called cyclin is needed for X to be active, and is degraded by Y.

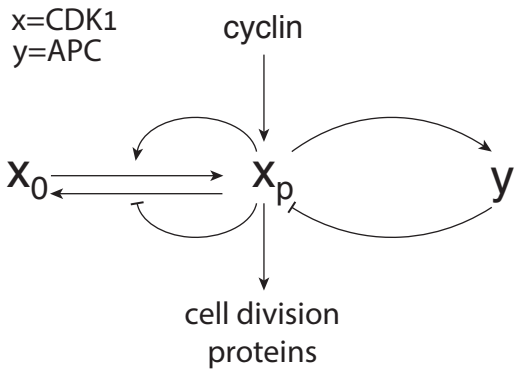

Figure 11.14

The positive feedback loop is due to the fact that Xp activates enzymes that increase its own phosphorylation- a form of positive autoregulation. A detail we will return to later is that Xp activates itself in two ways: activating the protein that phosphorylates x0 to make xp (cdc25), and inhibiting the opposing protein that dephosphorylates Xp to make Xo (wee1).

In the frog-egg cell-cycle circuit, the positive feedback loop causes bistability and **hysteresis**, as experimentally shown by Pomeraning, Sontag and Ferell (2003) and XXX. They used frog egg extracts and added a non-degradable version of cyclin in order to activate X – the more cyclin added, the more X is activated. When you start with zero cyclin and now slowly increase its amounts, Xp starts low and gradually rises (Fig 15). When cyclin reaches a certain threshold level, c_hi , the autoregulation kicks in and Xp jumps to a very high level, because it induces its own phosphorylation and inhibits its own dephosphorylation. Remarkably, when you start with high cyclin and now reduce its levels, Xp stays at a higher level then before, due to the autoregulation. Only when cyclin is lowered below a low threshold c_lo, does autoregulation become weak enough that Xp drops to its original level (Fig 15).

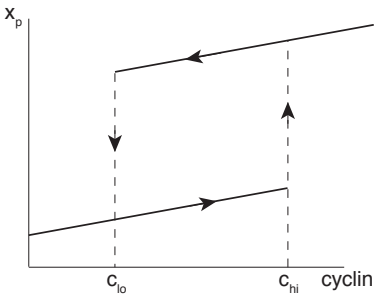

Figure 11.15

Here is how the oscillator works (Fig 16): at first y is low, and cyclin is transcribed at a constant rate, so that cyclin levels accumulate, and with them xp. When cyclin c reaches c_hi, xp shoots up. The clock goes tic. As a result Y rises, degrading cyclin. But thanks to the hysteresis, Xp goes down on the high arm, until cyclin drops below c_lo. Then xp crashes down and y drops, reseting the cycle. The clock goes toc. cyclin is no longer degraded by y and starts accumulating again, to begin another cycle.

The gap between the two transition points, c-hi and c-low, makes the transition from high xp to low xp robust to noise in the dynamics: A simple threshold mechanism would have just one threshold for the up and down transitions, making xp jitter up and down if cyclin dawdles around the threshold.

The levels of Xp show a slow increase and then an explosive spike. These crisp oscillations are characteristic of **relaxation oscillators**. The name alludes to the sudden relaxation of the tension built up as cyclin increases. The oscillation usually has an asymmetric pulse shapes with slow buildup at first, then accelerated buildup as the positive feedback kicks in, and a rapid decline, instead of the more symmetric pulses of delay oscillators or noise-induced oscillators (Fig 16).

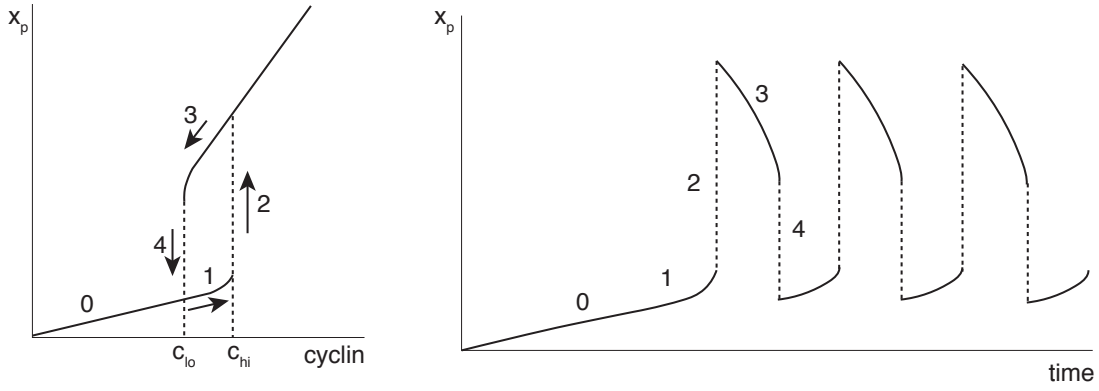

Figure 11.16

To change the frequency of this relaxation oscillator, all you need to do is change the rate at which cyclin accumulates (eg its transcription rate) (fig 17). The amplitude of the spikes remains almost unchanged. Easily tunable frequency is an advantage for a cell-cycle clock, because cell-cycle times range between 20 min in rapid embryonic development to days/weeks in adult tissues. Tunable frequency also occurs in heart cells, as our heart rate changes over a 2-3 fold range to meet our need for oxygen. Tunable frequency is harder to achieve in a simple delay oscillator without positive feedback, because changes to frequency are coupled to changes in amplitude.

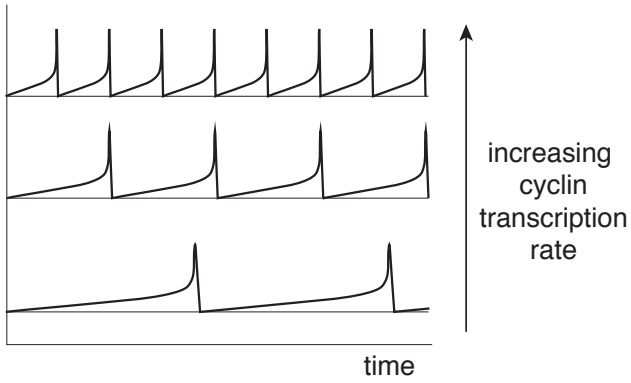

Figure 11.17

Robust bistability using two opposing positive feedback loops

To complete our look at the cell cycle oscillator, let's look in more detail at the positive feedback loop in circuit. As mentioned above, xp increases its own level in two ways, by activating the protein that phosphorylates xo to make xp, and by inhibiting the opposing protein that dephosphorylates xp back to xo. This makes two positive feedback loops.

Why two loops and not only one? I love the elegant answer proposed by James Ferell. The two-loop circuit can make bistability robust to wide variations in parameters.

To see this, we will use the rate plot method. Let's begin with no feedback, just production of Xp from Xo by phosphorylation, and the removal of Xp back to Xo by dephsophorylation (Fig 18). The rate of removal is a rising line, Xp times the rate of the phosphatase. The rate of production is the rate of the kinase times Xo. This is a decreasing line, that falls to zero when all of X is in the Xp form (Xp=Xtot), because there is no more Xo to be phosphorylated. The important part is where the two lines cross- crossing points are the steady-state points at which production equals removal. There is one crossing point, making a single steady state, with no bistability.

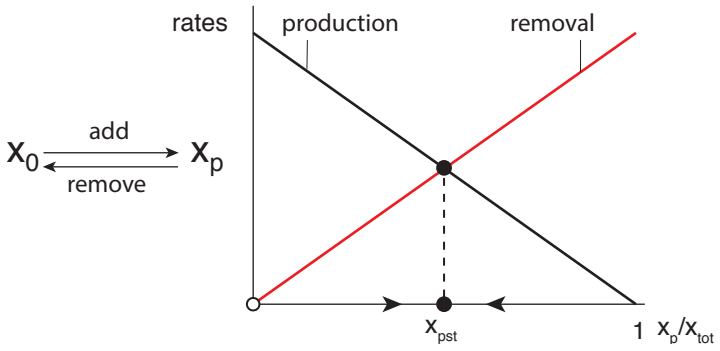

Figure 11.18

Now lets add one positive loop , in which Xp only activates its own production. This loop can show bistability, but this bistability is fragile. It is lost upon slight changes in the removal rate parameter. To see this, notice that feedback makes the production rate curve have a hump

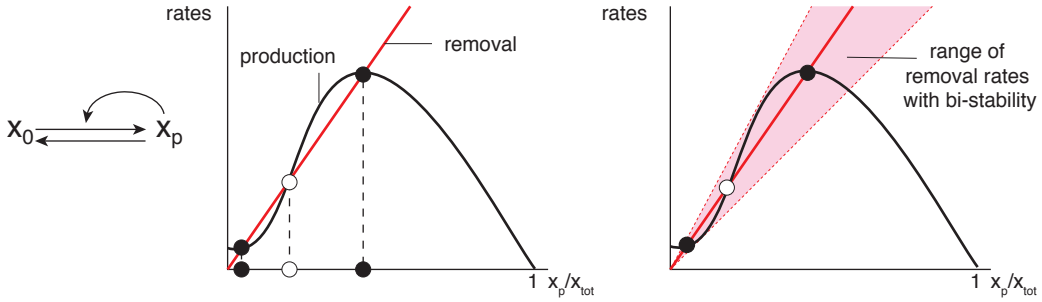

Figure 11.19

shape (Fig 19)- it rises with Xp due to the positive feedback (Xp activates its own production), and falls to zero when all of X is phosphorylated (Xp=Xtot). You can get bistability with three crossing points (a low and high steady state, and an intermediate unstable fixed point). However, it's enough that the removal rate line shifts slightly, due to a change in the number of phosphatases for example, and the three-fixed-point structure is lost. Bistabiltiy is not robust.

In contrast, the two-positive-loop design means that Xp both increases its own production and decreases its own removal. This creates a symmetry in the production and removal curves. Now both are hump-shaped (Fig 20). This makes their intersection points much less sensitive to

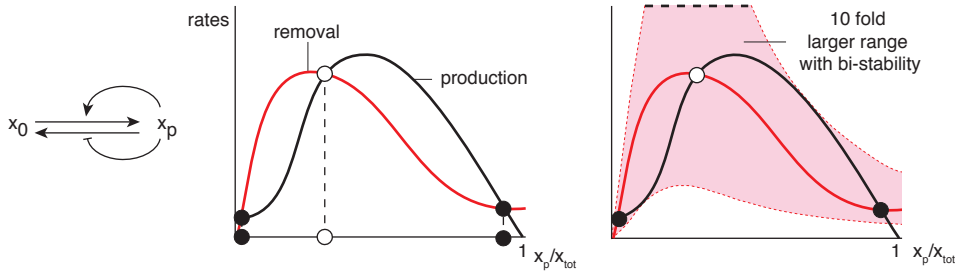

Figure 11.20

changes in parameters. The three-fixed-point structure survives changes in parameters that are 10-fold larger than the one-loop design, providing robust bistability and hence robust relaxation oscillations. Indeed, the two-positive-feedback loops are conserved in cell-cycle circuits throughout evolution from yeast to humans, highlighting their importance.

To sum up, oscillations require negative feedback and are aided by delays and cooperativity. Noise can turn damped oscillations into sustained pulsations with noisy amplitude and relatively precise frequency. Many biological oscillators, such as the cell cycle clock, use a motif in which a negative feedback loop is coupled to a positive feedback loop, resulting in spike-like pulses with tunable frequency and robust amplitude.