# The Complexity of Living Things

## Complexes and pathways

Although the basic mechanisms that underlie cellular biology are surprisingly few, there are many instances and many variations on these mechanisms, leading to an ocean of detail concerning (for instance) how the process of microtubule attachment to a centrosome differs across different species. Cellular-level systems, because they are so small, are also difficult to observe directly, which means that obtaining this detail experimentally is a long and arduous process, often involving tying together many pieces of indirect evidence. Most importantly, cellular biology is hard to understand because living things are extremely complex—in several different respects.

One source of complexity is the sheer number of objects that exist in a cell. At the molecular level of detail, there are thousands of different proteins in even the simplest one-celled organisms. These individual proteins can themselves be quite large, and assemblies of multiple proteins (appropriately called **protein complexes**) can be extremely intricate. One notable example for bacteria is the "molecular motor" which spins the **flagellum**—an assembly of dozens of copies of some twenty distinct proteins that functions as a highly efficient rotary motor. (See Figure 9.) This motor is atypical in some ways—most protein complexes are less well-understood, and do not resemble familiar mechanical devices like turbines—but it is far from unrivaled in its size or in the number of protein components. (Ribosomes, for instance, are much larger.) Unraveling this type of complexity is part of the discipline of biochemistry.

A **flagellum** is a whip-like appendage that certain bacteria have. It functions as a sort of propeller to help them move. An *E.coli* flagellum rotates at 100Hz, allowing the *E.coli* to cover 35 times its own diameter in a second.

A second type of complexity associated with living things are the complex ways in which proteins interact with each other, with the environment, and with the "central dogma" processes that lead to the production of other proteins. A *simplified* illustration of one of the best-studied such processes is shown in Figure 10, which illustrates how *E. coli* "turns on" the genes that are necessary to import lactose when

its preferred nutrient, glucose, is not present. Briefly, the gene *lacZ* is regulated by two proteins (called *CAP* and the *lac repressor protein*), which function by binding to the DNA near the site of the *lacZ* gene, and a feedback loop involving lactose and glucose affect the relative quantities of CAP and the lac repressor protein; however, as the figure shows, the details of this feedback process are nontrivial.

Many cell processes involve this sort of "interaction complexity," and often the interactions are far from being completely deciphered, let

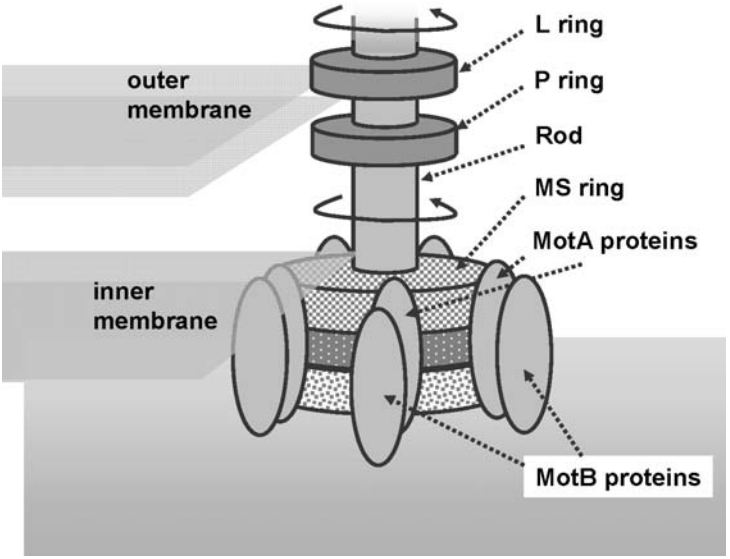


Structure of a bacterial flagellum (simplified). About 40 different proteins form this complex. The MS ring is made up of about 30 FliG subunits, and about 11 MotA/MotB protein pairs surround the MS ring. It is believed that these pairs, together with FliG, form an ion channel. As ions pass through the channel, conformational changes cause the MS ring to rotate, much like a waterwheel.

A similar "molecular motor" is used in ATP synthesis in a mitochondrion: rotation, driven by ions flowing through a channel, is the energy used to convert ADP to ATP. (See the section below, "Energy and Pathways").

**Figure 9. The bacterial flagellum.**

alone understood. Like the molecular motor that drives the flagellum, the chemical interactions in a cell have been optimized over billions of years of evolution, and like any highly-optimized process, they are extremely difficult to comprehend.

## Individual interactions can be complicated

Networks of chemical interactions like the one shown in Figure 10 are also complex in a different respect: not only is there a complex network that defines the *qualitative* interactions that take place, the

The *lacZ* gene is transcribed only when CAP binds to the CAP binding site, and when the *lac* repressor protein does not bind to the *lac* operator site.

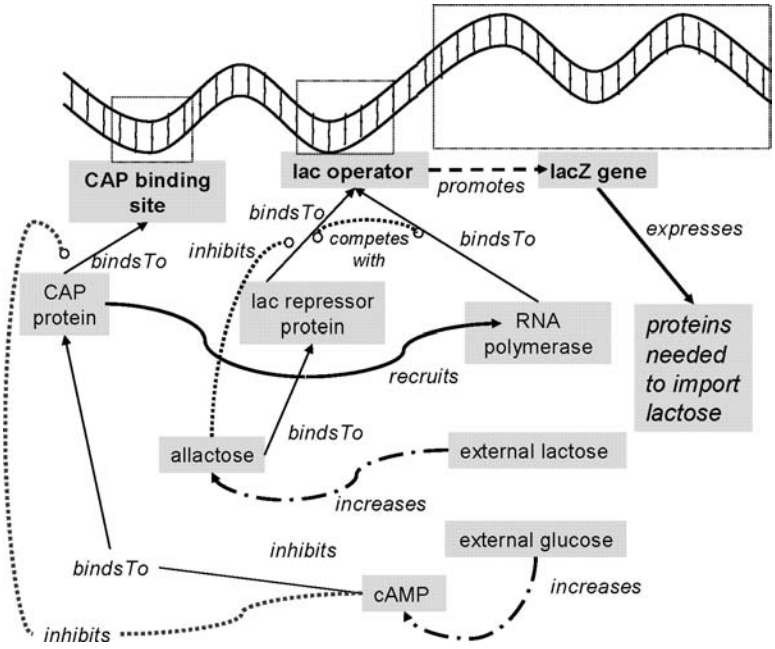


This network presents simplified view of why *E.coli* produces lactose-importing proteins only when lactose is present, and glucose is not.

**Figure 10. How *E. coli* responds to nutrients.**

individual interactions can be *quantitatively* complex. To take an example, increases in glucose *might* increase the quantity of cAMP linearly—but often there will be complex non-linear relationships between the parts of a biological chemical pathway.

The reason for this is that most biological reactions are mediated by **enzymes**—proteins that encourage a chemical change, without participating in that change. Figure 11 gives a "cartoon" illustrating how an enzyme might encourage or **catalyze** a simple change, in which molecule *S* is modified to form a new molecule *P*. It is also common for enzymes to catalyze reactions in which two molecules *S* and *T* combine to form a new product.

Enzymes can accelerate the rate of a chemical reaction by up to three orders of magnitude, so it is not a bad approximation to assume that a change (like S → P above) can only occur when an enzyme E is present. This means that if you assume a fixed amount of enzyme E and plot the rate of the chemical reaction (let's call this "velocity," V) against the amount of the substrate S (and like chemists, let's write the amount of S as [S]), the result will be the curve shown below. Velocity V will increase until the enzyme molecules are all being used at maximum speed, and then flatten out, as shown in Figure 12.

This model is due to Michaelis and Menten and is called "saturation kinetics." In fact, the shape of the curve shown is quite easy to derive from basic probability and a few additional assumptions—the ambitious reader can look at the mathematics in Figure 13 and Figure 14 to see this.

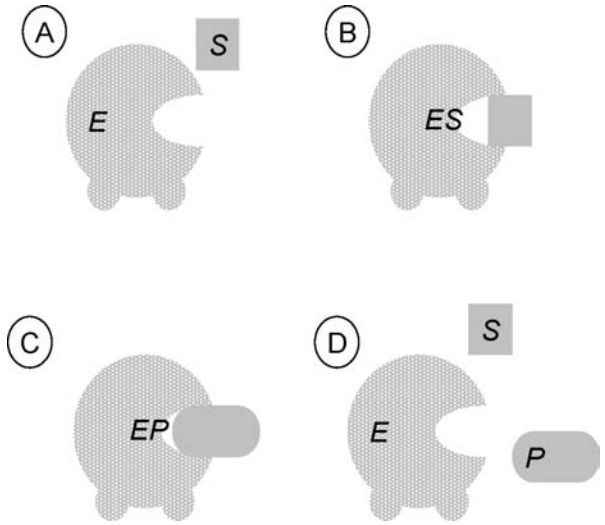


A cartoon showing how an enzyme catalyzes a change from *S* to *P*. (A) Initially, the enzyme *E* and "substrate" *S* are separate. (B) They then collide, and bind to form a "complex" *ES*. (C) While bound to *E*, forces on the substrate *S* cause it to change to form the "product" *P*. (D). The product is released, and the enzyme is ready to interact with another substrate molecule *S*. A chemist would summarize this as: E+S→ ES → EP → E+P

**Figure 11. How enzymes work.**

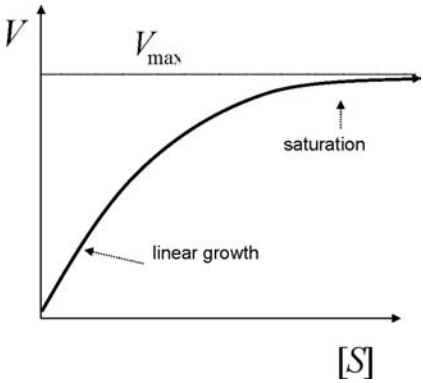


Reaction velocity with a fixed quantity of an enzyme *E*, and varying amounts of substrate *S*. When little substrate is present, an enzyme *E* to catalyze the reaction is quickly found, so reaction velocity *V* grows linearly in substrate quantity [*S*]. For large amounts of substrate, availability of enzymes *E* becomes a bottleneck.

**Figure 12. Saturation kinetics for enzymes.**

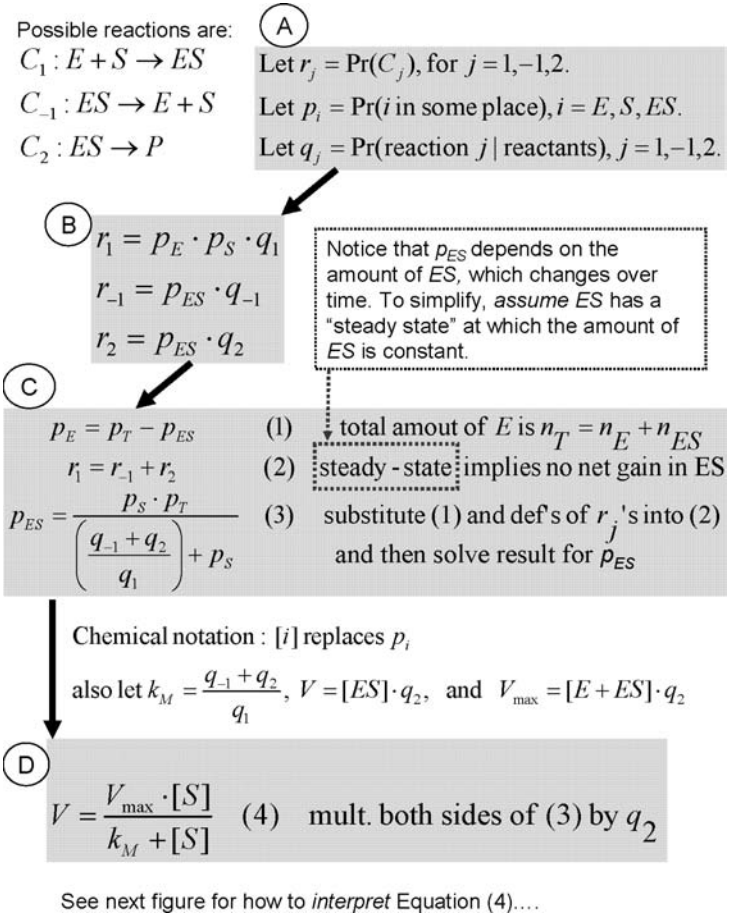


**Figure 13. Derivation of Michaelis-Menten saturation kinetics.**

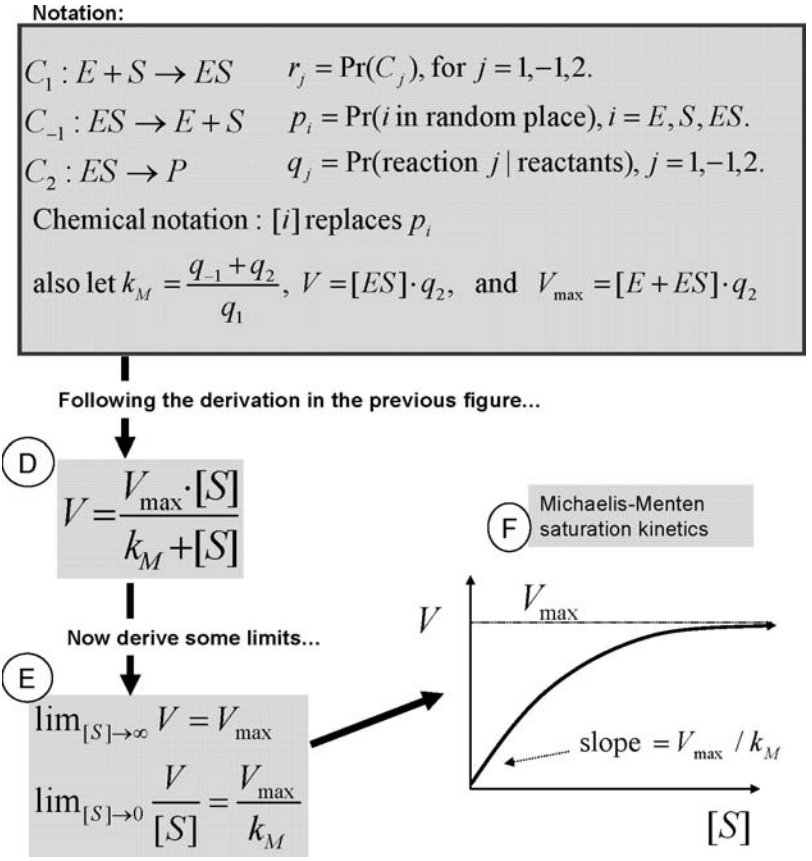


**The first limit shows that *V*, the velocity at which *P* is produced, will asymptote at $V_{max}$.**

**The second limit shows that for small concentrations of S, the velocity V will grow linearly with [*S*], at a rate of $V_{max}/k_{M.}$**

**Figure 14. Interpreting Michaelis-Menten saturation kinetics.**

Enzymes with more complicated structures can lead to more complicated velocity-concentration curves, as shown in Figure 15. A typical example would be an enzyme with two parts, each of which has an **active site** (a location at which the substrate *S* can bind), and each of which has two possible **conformations** or shapes. One conformation is a fast-binding shape, which has a high maximum velocity $V_{maxFast}$, and the other is a slower-binding shape with maximum velocity $V_{maxSlow}$. The lower part of the figure shows a simple state diagram, in which: (a) both parts of the enzyme change conformation at the same time, (b) shifts from the slow to fast conformation happen more frequently when the enzyme is binding the substrate, and (c) shifts from fast to slow tend to happen when the enzyme is "empty," i.e., not binding any substrate molecule. In this case, as substrate concentration increases, the enzymes in a solution will gradually shift conformation from slow-binding to fast-binding states, and the actual velocity-concentration plot will gradually shift from one saturation curve to another, producing a **sigmoid (**i.e., S-shaped) curve—shown in the top of the figure. A sigmoid is a smooth approximation of a step-function, which means that enzymes can act to switch activities on quite quickly.

A molecule that is composed of two identical subunits is a **dimer**; three identical subunits compose a **trimer**; and *N* identical subunits compose a **polymer**. An enzyme in which binding sites do not behave independently is an **allosteric** enzyme; in the example here, the enzyme exhibits **cooperative binding**.

Sigmoid curves and network structures are also familiar in computer science, and especially in machine learning: they are commonly used to define **neural networks**. A neural network is simply a directed graph in which the "activation level" of each node is a sigmoid function of the sum of the activation levels of all its input (i.e., parent) nodes. It is well-known that neural networks are very expressive computationally: for instance, finite-depth neural networks can compute any continuous function, and also any Boolean function. Although I am not familiar with any formal results showing this, it seems quite likely that protein-protein interaction networks governed by enzymatic reactions are also computationally expressive—most likely Turing-complete, in the case of feedback loops. This is another source of complexity in the study of living things.

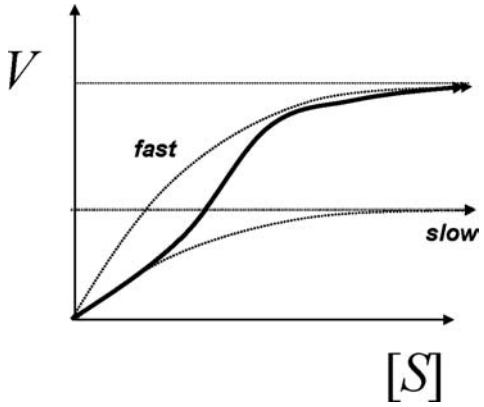


**Allosteric** enzymes switch from a slow-binding state to a fast-binding state, and tend to remain in the fast-binding state when the substrate *S* is common. Their kinetics follows a sigmoid curve.

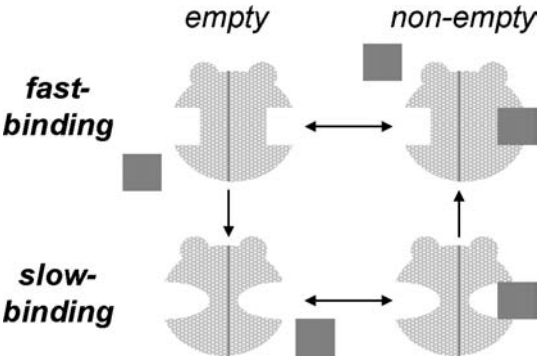


A typical allosteric enzyme: when one half is being used, the whole molecule tends to shift to the fast-binding state.

**Figure 15. An enzyme with a sigmoidal concentration-velocity curve.**

## Energy and pathways

Enzymes are important in another way. Running the machinery of the cell requires energy. Most of this energy is stored by pushing certain molecules into a high-energy state. The most common of these "fuel" molecules is **adenosine**, which can be found in two forms in the cell: **adenosine triphosphate (ATP)**, the higher-energy form, and **adenosine diphosphate (ADP)**, the lower-energy form. Enzymes are the means by which this energy is harnessed. Usually this is done by coupling some reaction P→Q that *requires* energy with a reaction like ATP→ADP, which releases energy. If you visualize the potential energy in a molecule as vertical position, you might think of this sort of enzyme as a sort of see-saw, in which one molecule's energy is increased, and another's is decreased, as in the figure below. (Dotted lines around a shape indicate a high-energy form of a molecule.)

More properly, ATP is combined with water to produce ADP plus inorganic phosphate, yielding energy: $ATP+H_20 \rightarrow ADP + Pi$. This reaction is called **hydrolysis**.

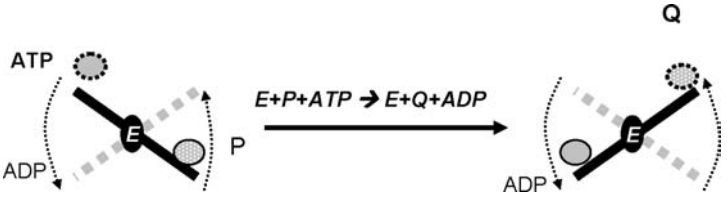


**Figure 16. A coupled reaction.**

Cellular operations that require or produce energy will often use an **enzymatic pathway—**a sequence of enzyme-catalyzed reactions, in which the output of one step becomes the input of the next. One well-known example of such a pathway is the TCA cycle, which is part of the machinery by which oxygen and sugar is converted into energy and carbon dioxide. A small part of this pathway is shown below in Figure 17. (Notice that this particular pathway produces energy, rather than consuming energy).

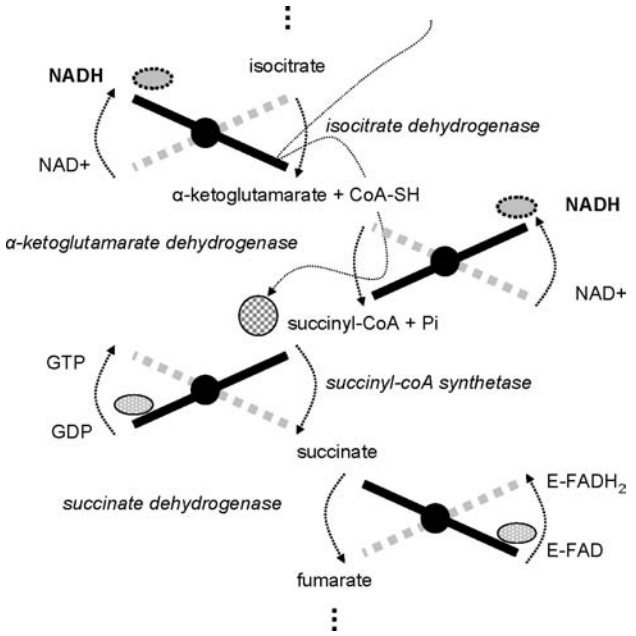


Part of the TCA cycle (also called the citric acid cycle or the Krebs cycle) in action. A high-energy molecule of isocitrate has been converted to a lower-energy molecule called α-ketoglutamarate and then to a still lower-energy molecule, succinyl-CoA (as shown by the path taken by the hashed circle). In the process two low-energy NAD+ molecules have been converted to high-energy NADH molecules. Each "see-saw" is an enzyme (named in italics) that couples the two reactions. The next steps in the cycle will convert the succinyl-CoA to succinate and then fumarate, producing two more high-energy molecules, GTP and $E\text{-}FADH_2$.

**Figure 17. Part of an energy-producing pathway.**

Since each intermediate chemical in the pathway (e.g., fumarate, succinate, etc.) is different, each enzyme is also different: thus a pathway that either consumes or produces large amounts of energy will often involve many different enzymes, again contributing to complexity.

## Amplification and pathways

Sometimes a pathway will act to amplify a weak initial signal. A good example of this is the pathway associated with **rhodopsin**. Rhodopsin is a G-linked protein receptor that detects light. Each rhodopsin protein cradles a "chromophore" molecule called **11-cis-retinal**. When a photon is absorbed by the 11-cis-retinal molecule, it changes shape, which causes rhodopsin to change shape and become "active." "Active" rhodopsin can then "activate" a second protein called **transducin**. Transducin, in turn, "activates" a third protein called **cGMP phosphodiesterase (PDE)**, an enzyme that hydrolyses a somewhat ATP-like molecule called **cyclic guanine monophosphate** (**cGMP**). In the **rod** and **cone** cells in the retina—the cells which sense light—cGMP acts somewhat like a chemical doorstop, propping open certain ion channels. When the concentration of cGMP is reduced, these ion channels close, changing the electrical charge of the cell and finally leading to a voltage signal. The process is thus something like this, where **R** is rhodopsin, **T** is transducin, and a***** denotes the active form:

The "fuel" used in a cell is chemically related to the bases of DNA and RNA. There are four **nucleobases** (aka **bases**) that form DNA: **adenosine**, **thymine**, **cytosine**, **and guanine**, abbreviated A, T, C, and G. (In RNA **uracil** replaces **thymine**.) A **nucleoside** is a base attached to a sugar: either **ribose** (for RNA) or **deoxyribose** (for DNA). A **nucleotide** is a nucleoside attached to a phosphate group: either mono-, di-, or triphosphate. These are abbreviated with 3- and 4-letter codes: e.g., ATP is adenosine triphosphate, and cAMP is cyclic adenosine monophosphate.

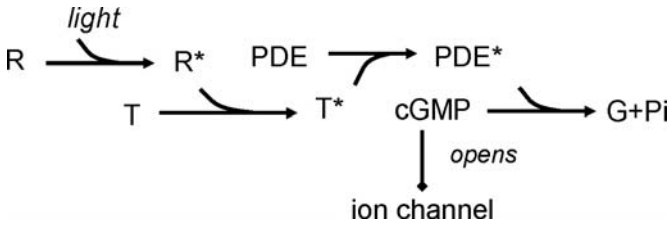


**Figure 18. How light is detected by rhodopsin.**

The interesting thing here, however, is that an active rhodopsin is unchanged after it activates a transducin, so it can go on and activate another transducin after the reaction completes. In fact, a single R* can activate thousands of transducin molecules per second, and likewise each PDE* can hydrolyze thousands of cGMPs per second. (A transducin can only activate one PDE, however.) This means that a single photon hitting the chromophore molecule can alter hundreds of thousands of cGMP molecules.

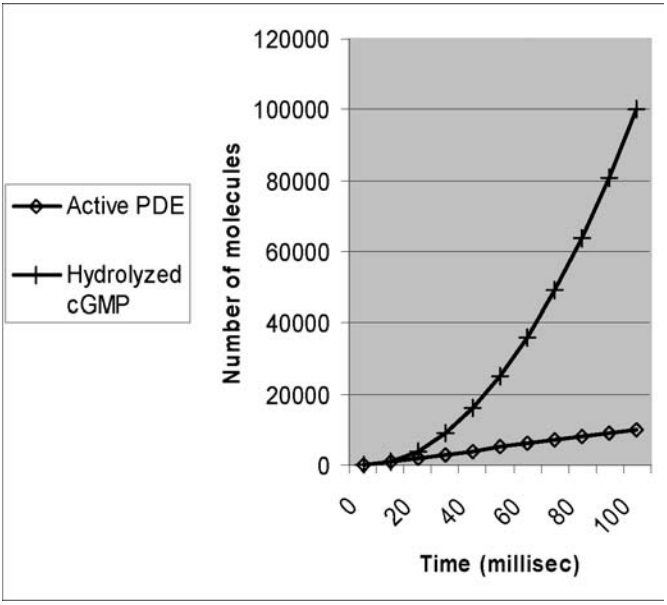


Number of molecules affected over time, assuming that each R* activates 100 transducin per second and each PDE hydrolyses 100 cGMP per second. (The actual numbers are larger).

The number of hydrolyzed cGMP molecules grows rapidly—at a quadratic rate—because it is product of two stages of "linear amplification". More stages of amplification would produce even steeper response curves.

**Figure 19. Amplification rates of two biological processes.**

In Figure 19, the pathway contains two "amplification" steps: both R* and PDE* affect more than one molecule each. Notice that the number of active transducin and PDE molecules grows linearly over time; however, since each PDE* hydrolyzes a linear number of cGMP molecules per unit time, the number of cGMP hydrolyzed grows quadratically over time.

## Modularity and locality in biology

Our understanding of macroscopic physical systems is guided by some simple principles—principles so universally applicable that we seldom think about them. One is the principle that *most effects are local*. This means that a good start to understanding how something works is to take it apart and see what touches what. Once we see that the ankle bone connects to the shin bone, we understand that those two components are likely to interact somehow.

This sort of common-sense approach to understanding systems fails for computer programs, where anything can affect anything. As a consequence, computer scientists are forced to construct elaborate schemes to limit the interactions of software components—in Java, for instance, private variables and methods, packages, and interfaces are all mechanisms for giving software constructs their own flavor of "locality." Programs that do not observe these principles are notoriously difficult to maintain, debug, and understand.

Like unconstrained software, the machinery of the cell also lacks "locality." A bacterium, for instance, is a complex machine, with thousands of *types* of parts (the types of gene products) and millions of *instances* of these parts. Although some of these parts form large structures (like the flagella), many of them are essentially just suspended in the fluid inside the cell. Components of the cellular machinery find each other, interact, and then separate, often without preference for a particular location.

This sort of non-local interaction is possible only for very small objects, and at very small scales. In a bacterium, proteins move about by **diffusion**, or random movement. In general, molecules around room temperature move very fast: for instance, a molecule of air moves at around 1000 miles per hour. However, molecules move

randomly, not systematically, which limits the ground that they cover. It is fairly easy to show that for objects moving by a random walk—specifically, objects that move a fixed distance in a random direction at each time step—the time it takes to cover a distance $x$ with high probability varies as $Vx^2$, where $V$ depends on distance traveled per unit time. This is very different from the macroscopic world, where the time to cover distance $x$ is usually linear in $x$.

The result of this is that diffusion is a very quick way of moving around for very short distances—say, the width of a bacterium—and a very slow way of moving around over larger distances—say, from the bar to the buffet table. This may be why very little internal structure is necessary for bacteria, or for the bacteria-sized organelles in eukaryotic cells: there is simply no need for it, since everything is already close enough to interact quickly with everything else.

An **organelle** is a discrete component of a cell. Some but not all organelles are membrane-enclosed areas.

Over objects as large as a typical eukaryotic cell, however, simple diffusion is not necessarily the most efficient way for molecules to find each other and interact. For instance, the enzymes used by cells to digest sugar are all localized to the inner membrane of the mitochondria—they still move by diffusion, but in a limited, two-dimensional area.[1] The various membranes and organelles in eukaryotic cells, therefore, do not only *limit* the way that proteins interact, by isolating some proteins from others—they also may improve the speed at which interactions within that enclosure take place, by limiting diffusion to a small area.

Besides diffusion, eukaryotes have a number of other mechanisms for **transport**: for instance, **vesicles** are small organelles that move

[1] Very approximately, cell membranes are about the viscosity of butter, while the cytoplasm of a cell is about as viscous as water, so molecules move about 100 times as slowly when they are stuck in a membrane. However, diffusion in two dimensions is asymptotically more efficient than in three dimensions, so it is faster to diffuse inside a membrane if the distance is large enough. Analysis of simple model systems suggests that the "cross-over point" at which membrane-bound diffusion is faster than simple diffusion is somewhere between the size of a bacterium and a mammalian cell.

materials from organelle to organelle; and within the cytoplasm, some proteins are hauled from place to place along microtubules, which are long fibers that run radially from the center of the cell to the periphery. Transport in eukaryotic cells leads to locality, and hence to some degree of modularity, which can be used to help understand cellular processes. More generally, the **subcellular location** at which proteins are found is often an important indicator of function.

It should also be emphasized that, while membranes provide some notion of locality inside a cell, membranes allow small molecules to diffuse through them, and biological membranes also have numerous mechanisms to allow (or actively encourage) certain larger molecules to pass through. Furthermore, because of properties similar to the random-walk property of diffusion, molecules that come close to an organelle tend to remain close to it for a while, and brush against it many times—Figure 20 gives some intuitions as to why this is true.

The result of this is that if receptors for a protein *p* cover even a small fraction of the surface of an organelle, the organelle will be surprisingly efficient at recognizing *p*. As an example, if only 0.02% of a typical eukaryotic cell's surface has a receptor for *p*, the cell will be about half as efficient as if the *entire* surface were coated with receptors for *p*. Cell-sized objects thus have a "high bandwidth"—they can recognize or absorb hundreds of different chemical signals, even if they are bounded by membranes.

To summarize, understanding even the "simplest" living organisms is far from simple. Analysis of how the different components of complex biological systems relate to one another is usually called **systems biology**.

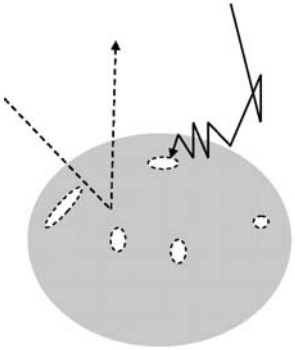

It can be shown that if a particle is released at distance $\delta$ from the surface of a sphere of radius $R$, it will touch the sphere before diffusing away with probability $p = R/(R+\delta)$. (See the book by Berg, 1983, cited in the last section, equations 3.1-3.5.) If the particle hits the sphere, bounces off, and returns to distance $\delta$ again, it has *another* chance to hit the sphere, again with probability $p$, so the expected number of times $n$ it hits the sphere before diffusing away is

$$E[n] = \sum_{n=0}^{0} n \cdot \Pr(\text{exactly } n \text{ hits})$$

$$= \sum_{n=0}^{0} n \cdot p^n (1-p) = \frac{p}{1-p} = \frac{R}{\delta}$$

This means that a protein nearing a relatively large membrane-enclosed object (like a cell or organelle) is more likely to follow a path like the solid line than the dashed line—it will typically hit the cell many times before diffusing away, giving it many chances to "find" a receptor.

**Figure 20. Behavior of particles moving by diffusion.**