

I. Conservation laws for isolated systems

Many of the most fundamental laws of science are *conservation laws*. In simplest form, a conservation law says that in an isolated system the total amount of some kind of stuff remains the same. What the ‘stuff’ is depends on the system and its properties.

What ‘isolated’ means also varies. Ideally, being isolated means that a system has no interaction whatever with the rest of the universe. In practice, a system need not be totally isolated in order to be practically isolated in some respects. The inside of a thermos bottle is very close to being thermally isolated, even though it is not at all mechanically isolated (i.e. you can pick it up and carry it around without transferring heat into or out of the stuff inside). An insulated copper wire is electrically isolated even though it may not be thermally isolated.

Classical physics is based on three fundamental facts that apply to almost all systems:

- (1) Conservation of Mass: under normal circumstances, the total mass of an isolated system is constant. This is just the statement that matter does not vanish into nothingness or appear out of nothingness, and it is such a common-sense principle that I have no idea when scientists first decided to start writing it down.
- (2) Conservation of Momentum: the total momentum in an isolated system is constant. Galileo recognized this, and Newton enunciated it in his First Law of Motion: “An object in motion tends to remain in motion unless acted upon by an outside force”. Angular momentum is also conserved.
- (3) Conservation of Energy: under normal circumstances, the total amount of energy in an isolated system is constant. This law has a colorful and exciting history of which I know very little. It was probably enunciated soon after Newton’s laws, but was argued about for over a hundred years. The sticky point was probably what happens to energy when friction operates, and it was not until the 1800’s that people could measure heat accurately that the equivalence of mechanical and heat energy could be established.

The reason for the caveat “under normal circumstances” in the mass and energy conservation laws is that these laws are not perfectly true; they are only approximations that hold extremely accurately for systems that we are likely to encounter on earth. These approximations are actually special cases of

- (4) Einstein’s relativistic mass-energy conservation law: the total energy plus c^2 times the total mass is constant for any isolated system, where c is the speed of light.

Laws (2) and (4) are still exact laws of nature as far as anyone knows.

Modern physics has added a couple of new laws - here the term “modern” means within the last two hundred years or so. One of the first discoveries in electricity was that if you built up a static charge on something and treated it carefully, the charge would stay for a long time. As this observation was checked more and more carefully over the years, it turned out to be (and remains) true as accurately as anyone can tell, and it is known as

(5) Conservation of Charge: the total electric charge in a completely isolated system is constant. This holds even if you have electrons and protons undergoing nuclear reactions to form neutrons (and vice versa), or other subatomic processes.

There is another conservation law, I think, that is important to astrophysicists and nuclear physicists:

(6) Conservation of Baryon Number: the total number of baryons (protons and neutrons) is constant, even if they undergo nuclear reactions and turn into each other. (I might easily have stated this incorrectly, by the way. Ask a real physicist. It is not important in Math 254, thanks goodness.)

Often some stuff that is not perfectly conserved may be almost conserved to an extremely good approximation. We saw an example of this already: the error in mass and/or energy conservation is tiny when all velocities are much smaller than the speed of light. Another example comes from chemistry. When you burn a piece of wood, you start with a certain amount of carbon and a certain amount of oxygen, and you end up with a certain amount of carbon dioxide. The wood is not conserved and neither is the amount of oxygen gas (i.e. oxygen in the form of O_2 molecules). But to a very good approximation the number of carbon atoms at the end of the process is the same as the number at the beginning, and the same applies to the oxygen atoms. This principle applies to almost all chemistry, and can be stated as:

(7) Conservation of Elements: in an isolated system, in which nuclear reactions may be ignored, the total amount of each chemical element is constant.

And, of course, if there are no chemical reactions either (or if they are negligibly slow) the same principle applies to all the substances in the system

(8) Conservation of Substances: in an isolated system, in which nuclear and chemical reactions may be ignored, the total amount of each substance is constant.

This principle applies to all kinds of mixing problems, e.g. salt in seawater, pollutants in rivers and groundwater, high-oxygen blood and low-oxygen blood in the circulatory system.

Other fields have other conservation laws, though the more complicated the phenomena

the smaller the role of the conservation law. In animal ecology, we might formulate the law of

(9) Conservation of Animals: in an isolated ecosystem, in which no births or deaths occur, the population remains constant.

Economics has the law of

(10) Conservation of money: in an isolated economy, in which neither the government nor individuals print money, the total amount of money remains constant.

As you see, some of the assumptions implicit in conservation laws may get pretty suspect!

Conservation laws for isolated systems are easy to write as differential equations. Let $Q(t)$ denote the quantity of conserved stuff in the system at time t . (It is a bit silly to write the t dependence of something that we have just said is going to be constant, but at least it's not wrong. Writing it this way will help us later, when we discuss nonisolated systems in which $Q(t)$ varies.) The fact that the stuff is conserved means $Q(t)$ is constant, which implies

$$\frac{dQ}{dt} = 0. \tag{1}$$

Thus, whenever we have an isolated system with a conservation law, we have at least one very simple differential equation.

II. Conservation laws for interacting components

The conservation laws we have just described are really idealizations that people have constructed on the basis of huge amounts of observation of real systems. By doing careful experiments, a scientist can approximate mechanically or thermally isolated systems and extrapolate their conclusions to the ideal case of perfectly isolated systems. But almost everyone (including most scientists) does not really care about isolated systems - we want to understand systems that have lots of interactions with lots of other systems!

Why is this? Isolated-system conservation laws are only really useful when applied to very simple systems. You can use momentum conservation to determine the motion of a thrown ball, because when you know its mass and momentum you know everything about its motion. But if you throw a complicated object like two balls connected by a spring, knowing the total momentum says nothing about how they are whirling around each other or how they are affecting each other via the spring. Similarly, we know that the Earth is pretty well isolated chemically, hence the amounts of carbon, silicon, oxygen, and other elements are constant. However, if we want to model the Greenhouse Effect or the Ozone Hole, we need a whole lot more information than just the overall element conservation laws!

Conservation laws can tell us lots about complicated systems with many components, but we have to develop a new form in which to state them. The fundamental idea is still that conserved stuff is not created or destroyed, but if the system has several components then the conserved stuff can move around amongst the components. The rates at which the conserved stuff moves between the components are called the *fluxes*. The conservation law no longer says that quantities of conserved stuff are constant. Instead:

The rate of change of the quantity of conserved stuff in a component of a system equals the total flux arriving from all the other components.

The system does not even need to be isolated any more. We can think of the outside universe as just another component, able to exchange stuff with the normal components.

To make this principle explicit, suppose we can divide the system into some number N of simple components. Here “simple” means that if we know how much conserved stuff the component has, we know everything about it. Let $Q_i(t)$ be the quantity of conserved stuff in component i at time t . Then

$$\frac{dQ_i}{dt} = \sum_{j \neq i} F_{j \rightarrow i}(t), \quad (2)$$

where $F_{j \rightarrow i}(t)$ denotes the flux *from* component j *to* component i at time t . These flux functions have a very important mathematical property:

$$F_{i \rightarrow j}(t) = -F_{j \rightarrow i}(t). \quad (3)$$

This is simply the statement that the rate at which i gains stuff from j equals the rate at which j loses stuff to i , with a loss rate considered equivalent to a negative gain rate.

We now have a way of writing the fundamental conservation laws of science that can actually be useful in understanding real complicated systems. However, we cannot say what the system is going to do because we don’t know the fluxes. So how can we find out what the fluxes are? We have to go and look in great detail at the specific system we are given and see what the components are, how they are connected, and so on. This will give us some very specific equations, sometimes quite complicated for the fluxes. These equations are called the *constitutive relations* of the system. When we have the constitutive relations as well as the conservation laws, we have a complete mathematical description of the system in terms of differential equations. The rest of the problem is then the purely mathematical task of finding solutions to the differential equations.

It's worth restating the conclusion of the last section: MANY real systems are described mathematically by a combination of *conservation laws* and *constitutive relations*.

The conservation laws are the easy part. They have simple interpretations, and they tend to be extremely good approximations to reality. They are very general; the same conservation laws apply to huge numbers of different systems. The conservation laws yield most of the time-derivative terms in the differential equations.

The constitutive relations are the hard part. They are specific to each specific system, and typically when you go to a new system you have to come up with new constitutive relations. They may be complicated, and they often involve substantial idealizations and approximations of reality. If there is a discrepancy between theory and reality, it is probably in a constitutive relation.

The plus side is that, since the fundamental laws are embodied in the conservation laws, it is not so essential to get exactly the right constitutive relation. Often quite a coarse approximation gives a very useful mathematical model. It frequently turns out that we can do very well with constitutive relations that are *linear functions* with *constant coefficients*. With this kind of approximation, the differential equations end up being linear with constant coefficients also, and we can then solve them using the techniques of Chapter 4. This is why linear DE's with constant coefficients are so important, although it seems at first sight like a very specialized branch of DE's in general.

Of course, there is an art to approximating difficult systems by linear DE's with constant coefficients. The name of this art is "Mathematical Modeling", and you will see plenty of it in this and your subsequent courses.

One final important remark. This discussion of 'conserved stuff' and 'fluxes amongst components' makes lots of sense when we are talking about substances like water and sugar, or chemical elements like hydrogen or silicon. It even makes sense with less-tangible entities like energy and entropy - you can sort of visualize energy going from kinetic energy in a waterfall to electricity in a hydroelectric power station, through transmission lines to a toaster or a TV, and emerging as heat or light or sound. But what about momentum? How does it flow?

The only way to change the momentum of an object, according to Newton's Second Law, is to exert a force on it. Therefore, the mysterious concept 'flux of momentum' is actually the familiar concept of force. Or, in symbols:

$$F_{j \rightarrow i}(t) = \text{force exerted on component } i \text{ by component } j \text{ at time } t. \quad (4)$$

The mathematical property (3) that all flux functions must satisfy now becomes Newton's Third Law: if one object exerts a force on another, then the second also exerts a force on the first that is equal and opposite to the force the first exerts on the second.