

Surviving The Cyber Age

Chapter 5 - Control

Table of Contents

- Chapter 5 - Control.....1
- Control mechanisms.....2
 - A simple control system.....2
 - Control equations.....3
- Environmental conditions.....6
- Natural control systems and evolution.....7
 - Flows between disparate temperatures, pressures, and so forth.....7
 - Natural control systems that fight against entropy.....8
- The hierarchy of operational capabilities.....11
 - Biological processing and movement rates.....11
 - Computer processing and movement rates.....13
 - A simple transistor control loop.....13
 - A digital switching system.....14
 - Control systems with combinational logic only.....15
 - Storage and sequential finite state machines.....15
 - Computer hardware.....16
 - Instruction Sets and General Purpose Computers.....17
- Larger control systems.....18
 - Timing and the Implications to physical structures.....18
 - Operating parameter limitations.....19
 - Interference limitations.....19
 - Multiple frequencies within composites.....20
 - Timing and coordination at distance.....20
 - Shielding.....21
 - Central vs Distributed systems.....22
 - Parallel processing.....22
 - Communications and control architecture.....23
 - Up the cognitive ladder.....24
- Fundamental control system limits.....24

Control mechanisms

Control systems turn sensor inputs into actuator outputs through communications channels.

A simple control system

One of the simplest control systems is the common toilet flushing system used in homes around the world. It consists of two related components; the flusher, and the filler.

- The toilet is a composite composed of a tank, a dump pipe, a seat with integrated drain pipe and water source, the flusher, and the filler.
 - The flusher is a composite composed of handle or button connected to a flapper where the handle lifts the flapper to dump the water from the tank into the toilet, and the flapper has an air-filled area that will float if fully underwater but not if partially closed, and seals to the dumping pipe when closed.
 - The filler is a composite composed of a float, a water flow valve, and an overflow drain. The float floats at the current level of the water in the tank because its density is less than water and more than air. It is attached to the valve so that the valve is open for flow of water from the water source when the float is below a preset level, and as the float rises, it closes the valve. The overflow drain dumps overflow of water into the toilet in case the float or valve fail to operate as intended.

In combination, these components provide a control system in which, in sequence:

- the user activates the flusher actuator (the flapper actuator through the handle),
- the water in the tank dumps into the seat thus flushing waste down the drain pipe,
- the float lowers to below the level where it actuates the flow valve to flow water,
- the water flow valve starts to fill the tank at a rate less than the dump pipe empties it,
- the tank gets empty enough for the flapper actuator to close and seal the dump pipe,
- the water fills the tank to the level where the flapper is sealed against the dump pipe,
- the water keeps flowing though the flow valve lifting the float sensor,
- the float sensor gets closer to the preset level actuating to close the flow valve,
- as the float approaches the preset level, the water flow valve shuts completely.

At this point, the composite is ready to go through the cycle again. However, as a safety precaution against a failure elsewhere in the system the overflow drain assures that incoming water will not end up overflowing the tank. As a control system:

- **Sensors:** The flush handle, the float, and the flapper flotation;
- **Actuators:** The water flow valve, the flush handle, and the flapper;
- **Communications:** The connections from the flush handle to the flapper, the float to flow valve, the flapper to the water and its hinge, and the water to the flapper and float;
- **Controls:** The flow valve mechanism, the flapper flotation and seal physics.

The composite is actually two different control systems that interact with each other; one is the flushing mechanism, and the other is the filler mechanism. They are composed in the environment of the toilet to form the larger control system. Without the flusher, the filler control system can operate independently, and something of the same sort is used to refill ponds and fountains with the trigger for starting the fill as the level of the water, which goes down with evaporation and splashing, causing the float to trigger the flow valve. Water tanks in water systems operate in much the same way, however, the fill level is, in many cases, controlled remotely, and multiple sensors are used to detect water levels, water pressure from outside the tank, and other parameters to fulfill other aspects of the overall operation in the context of an overall water system, resulting in a more complex but more capable system.

Control equations

Most control systems can be modeled by mathematical equations to allow the system to be designed for reliable operation across a range of operating conditions. The toilet flushing system just described will only work under specific gravitational, temperature, motion, and pressure conditions with forces applied within identifiable ranges. When it gets out of design parameters, the system will fail in various ways:

- If the water is too cold, it will freeze, if it is too hot, it will turn to steam, and either way, the toilet will not work as described.
- In low gravity, the flapper will not close at the same rate and may not seal, while in high gravity, the leverage required to lift the flapper may not be enough for the user to start the operation of the overall system.
- If the environment is bouncing up and down, like in an airplane under turbulence, the water will spill all over the place, the float will bounce up and down, and the flapper may not close in time or seal. Thus airplane toilets work differently.
- If the air pressure is too low, the float may rise higher making an incomplete tank fill, and if it is too high, it may force the flapper into the dump pipe, never to return. If the water pressure is too low, it may take too long to fill the tank for practical use, and if it is too high, it might burst or freeze the control valve causing the tank to never fill or always fill to overflow.

There are certainly other failure modes, such as the water filling too fast so that the system always overflows, thus wasting water, and the overflow may be too low, so the flow valve never shuts completely, again wasting water, and there are other conditions related to timing that might make the system unstable and never get to a rest state.

The engineering of a toilet takes all of these factors and others into account in creating a control system designed to operate under desired conditions for its use. Figuring out the properties of components so as to stay within the bounds of the system requirements is usually done by using various equations. For example, the size, flow rate, change rate, composition, and cost of the flow valve might be limited to:

- A size allowing it to fit in the tank while still allowing enough space for everything else,
- A flow rate that fills the tank in the desired time frame based on water pressure and tank size, (30 seconds for 6 liters $\rightarrow 6/30=0.2$ liters/second, but at 60s, $6/60=0.1$ l/s)

- Handling a change rate so that turning the valve on and off supports the fill rate timing requirements (if it takes 9 seconds to go from open to close and 1 second to go from closed to open, and if the average flow over these times is half the maximum flow rate, that means that 10 of those 30 seconds for a refill are operating at half the flow rate, so the average flow rate is 25/30 of the rate identified above. So we need to flow at 30/25ths of the previous equation, or 0.24 liters/second for a 30 second refill.
- Made of materials that can operate in the air and under water and last for however many years the design requires while performing a predetermined number of refills per hour and total number of refills with only limited leakage,
- At a cost so that the total cost of all the parts is less than a defined cost requirement.

The same sort of requirements will drive all the other components as well as the composites and their properties, and the combined properties of the composites interacting in the operating environment must allow it to operate properly over its designed operating range.

Classes of control systems and failure mechanisms

The toilet control system is essentially a '**bang-bang**' control system. It operates between two extremes, full and empty, and in that sense is similar to a digital system (full = inbound flow off, and empty = inbound flow on). Of course near the full and empty the system goes analog, slowing inbound flow as the valve closes and increasing it as the valve opens. Nothing is instantaneous and if we tried to drive it to that extreme, something would break or the system would explode. For example, to fill the whole tank in 0 seconds, the flow rate would have to be $6/0$ liters per second, and anything divided by 0 is infinite. We don't have pipes that can handle infinite volumes of water per unit time. The same is true for electronic systems, biological systems, nuclear systems, and all other systems.

Continuous feedback control systems allow things like flows to vary over a range so that, for example, instead of opening or closing a valve, a water tank feeding a municipality can refill all the time with the amount of water entering the tank per minute equal to the average amount of water leaving the tank over the last minute. Since the averages change slowly, the control valve will not rush between extremes, and the rate of opening and closing the valve can be reasonably slow, reducing stress on the materials, and so forth. This approach uses feedback to control the system.

Feed-forward control systems are designed to anticipate the future needs of the system so that by the time the system needs whatever the control system is controlling, it's already there. A clever system might determine that during half-time at a major football game, there will be more water flow required in the stadium than during the periods when the game is actively being played, and increase the flow of water into the system to fill the tanks to maximum levels before half-time rather than sizing the incoming flow to be able to meet the maximum demand all the time. But even a less clever system can determine demand before it exists if the engineering process has devised equations to predict the future demand.

Positive feedback is an example of how this can go wrong. Of course positive feedback in a social situation is often used to reward a speaker for doing a good job. But we are talking about a control system, and the feedback is in the sensor data collected from the media as a result of the actuator actions and their effects on the media.

A phenomena called '**ringing**' happens in most systems what a valve opens or closes. That's because, for example, the closure rate of a valve might be slower than the time it takes to overflow the system, resulting in an overflow that then has to be emptied. But the emptying mechanism might empty a bucket at a time, and the bucket might be bigger than the overflow amount, causing the system to become under-filled again. Then the valve has to start flowing again, and if it over-shoots we will need another emptying bucket, and so forth.

If the timing of these operations is right, the system settles down over time (damped), but if the timing is wrong (under-damped), then the emptying happens just as the filling has to turn back on, and the filling happens just as the emptying gets started, and the system ends up emptying and filling faster and faster until the whole system is just wasting water all the time. Worse yet, as the mechanical changes of pressure in the tank caused by the timing of these actions increases and decreases the pressure on the walls of the tank, it might cause the tank itself to start resonating with the flows, and the tank might start to break as well.

Waves are present in physical systems, and all systems are physical. In the water system, when you put water into the system, it increases the amount of water where the water entered, but nothing happens instantly in the physical world, so the water level will be increased only where the water was added. Water seeks it own level, so the water at higher levels will fill in at lower levels, and the filling in process goes at a finite rate. You get waves.

Science does not fully understand everything about waves in water yet, particularly in the oceans, but one thing we know is that waves 'add' and 'subtract' when they 'interact'. In water, two wave peaks arriving from two different directions will produce peaks and valleys more extreme than the individual waves produce on their own. That's because water is an incompressible fluid (you cannot squeeze it into a smaller space), so more water in the same place (a column horizontally confined because there isn't enough time to go anywhere else and confined at the bottom by the bottom of the space in a situation where gravity pulls toward the bottom) has to go up, and less water has to go down.

When the timing of the sources and drains of water in the space of the tank are such that: water is drained in places and at times where there are valleys, and water is filled at places and times where there are peaks, the peaks will get higher and the valleys lower. If this continues to happen over time, the waves will get bigger and bigger, putting more force on the walls of the tank, till the tank can no longer take it. This is sometimes called '**sympathetic vibration**'. Water filling and emptying is just one example. Wind on water increases wave sizes, waves bounce off of surfaces (reflection), and the same things happen with air pushing back and forth on structures, and people walking in step on surfaces, and wind coming into windows on cars, and so forth. Resonant frequencies (changes at frequencies that create positive feedback) produce increasing amplification of wave heights over time.

The design of the system should reasonably anticipate these conditions and close the valve sooner so it never quite overfills, make the buckets smaller, and so forth. But this can only go so far, because the components are not perfect and they change slightly with temperature, usage, and a wide range of other things over their useful period of operation.

All of these things have to do with the nature of time and the fact that control systems operate over time frames. Treating them as if they were static ignores the dimension of time, and can lead to disaster.

Examples of substantial real world effects from waveforms include things like singing in the shower sounding better (louder and more pure in tonality because of resonance with the shower) than you really sound to other people in a room, singers able to break a glass with their high notes (by singing at resonant frequencies with the glass causing it to vibrate its crystalline structure into structural collapse), and bridges falling down when soldiers march across them (the resonant frequencies of footfalls with the bridge causing the bridge to collapse).^{1 2} Troops are generally ordered to not do this any more. A really impressive example is the 1940 Tacoma Narrows Bridge collapse.³ In this case, the frequency of wind changes from the water body it crossed caused the suspension bridge to form waves that eventually destroyed the bridge surface, broke suspension cables, and bad things followed.

Environmental conditions

All control systems operate within environments, including the media between actuator and sensor, the overarching context of what is being controlled, the parameters bounding the envelope of desired behavior, and the larger context (e.g., a toilet on Venus or Saturn).

Environmental conditions for such systems can be natural or artificial, accidental or intentional, and friendly or malicious depending on your perspective. They can effect the control system mechanisms, the operating conditions, and the media.

- **Artificial accidental friendly** is the most common situation for human-created control systems, in that the systems are: put in place by **human design** to achieve specific desired controls over **an artificial system**, like a municipality clean water tank; operated in a **friendly** environment where the people **want the system to work** to keep providing water **reliably** and animals and other similar potential pests are blocked from entry and encouraged to leave; and things that violate the design constraints are **accidental** in that they were **not anticipated or commonplace** in the design environment and **not the result of intent to damage** the system.
 - One of the things that's important to understanding accidental damage is the concept of due diligence. Basically, it is the responsibility of a responsible designer to know the history of the sorts of things they design and to understand what can go wrong and has gone wrong in the past so as to anticipate the possible failures modes, their likelihood under the intended design environment, and to document decisions made by management to constrain the design after management has been informed of the potential consequences of these decisions.
- **Artificial intentional malicious** is a very different part of the space, and a part of the space faced by many parts of the Internet today. But let's keep it to simpler systems for now, like the control systems of a water system operating in an environment where there is a war or other similar conflict underway. The **artificial** part remains the same except of course the environment of the intended design is usually quite different. Intentional malicious attack against a water system includes things like **destroying the electrical supply** to the pumps and control mechanisms over the valves, **firing weapons** of various sorts **at the tank and water supply**, **poisoning** the water with

1 https://en.wikipedia.org/wiki/Angers_Bridge April 16, 1850 collapse

2 <https://www.londonremembers.com/memorials/albert-bridge-troops>

3 [https://en.wikipedia.org/wiki/Tacoma_Narrows_Bridge_\(1940\)](https://en.wikipedia.org/wiki/Tacoma_Narrows_Bridge_(1940))

things that cause the **valve to stick**, the **pipes to get gummed up**, and so forth. Nothing in the specification dealt with the usability of the water itself, so poisoning the water in terms of output safety is not an issue for this control system. Note that water can be potable (usable in a pot for cooking) or not, and for non-potable water, the requirements of chemical composition and infestations are different than for potable human supply. The **unfriendly** nature of the **intentional** component means that we cannot trust people to want the system to keep working, and that implies that the facility itself must be protected from intrusions and alterations.

- Diligence for systems intended for this use are different in that the historical information should include understanding of weapons and weapon effects, but the specification likely does not cover direct nuclear explosions at the water tank itself. However, insider threat actors likely should be anticipated in the unfriendly environment, and explosions not directly targeting and hitting the tank should be considered in terms of the structural components used and the facility within which the water tank is placed. Perhaps it should be under ground? An artificial lake in a valley so that the water stays there even if there are major explosions? And so forth. Redundant supply, drains, and control systems might be reasonable, and detection and response times for fixing damage might be part of the overarching control system design.

The full set of alternatives is obviously a much larger set than this, but you should, by now, get the idea. The environment leads to the design of an artificial system and its components and composites.

Natural control systems and evolution

The simple artificial control system discussed above is indeed simple compared to most natural control systems. So let's start with the basics...

Flows between disparate temperatures, pressures, and so forth

The naturally occurring self-adjusting control systems of weather conditions are based on the physics of thermodynamics. "Thermo" as in temperature, "dynamics" as in what happens over time. Generally, if there is something hot next to something cold, the cold thing will get hotter and the hot thing cooler because heat transfers between them at a rate depending on the particulars of the situation. If you pour hot water into a bathtub full of cold water, the water in the tub will heat up, but not instantly, and not to the temperature of the incoming hot water. Ignoring the temperature of the tub itself and the surrounding atmosphere, the temperature will eventually reach the average of the volume times temperature of the hot and cold water portions. If 10% (0.1) of the water was at 150 degrees and 90% (0.9) was at 50 degrees, the average will be $150 \cdot 0.1 + 50 \cdot 0.9 = 15 + 45 = 60$ degrees. How long it takes to get to where the water has approximately that temperature everywhere depends on properties of the water and how much mixing there was. If there is a mixer in the water (and assuming the mixing doesn't change the eventual outcome, which it normally will), the average will be reached throughout the tub sooner than without the mixing. The 'heat equation' details how this works.⁴

⁴ https://en.wikipedia.org/wiki/Heat_equation provides some details

In the air, hotter air mixes with colder air producing wind as the temperature of air leads to different pressures. While water in an incompressible fluid, air is compressible gas, and the equation $PV=nRT$ tells us that if the temperature (T) is higher, all other things being equal, the pressure (P) will be higher. Air flows from high pressure to low pressure, which is to say there will be wind as the higher pressure air flows into the lower pressure air to equalize pressure. Of course cold air is also heavier, as we all know that 'hot air rises', which is really to say, 'the cold air pushes it up'. The natural tendency for differences to become non-differences through physics is called '**entropy**', often thought of as randomness. Entropy is a measure of the randomness. or the degree to which energy has dispersed within a system

Thermodynamics has 3 'laws', which is to say, things that must always be true according to physics:

- The first law states energy cannot be created or destroyed, only transformed.
- The second law states that the total **entropy** of an isolated system can only increase over time.
- The third law states that the entropy of a closed system at thermodynamic equilibrium approaches a constant value when its temperature approaches absolute zero

These various natural mechanisms for motion of temperature and material flow as well as others are not control systems, in that there are no sensor, actuator, communication, and control mechanisms involved. But they are feedback systems in which a change such as adding hot water to a cold water tub cools the hotter water and heats the cooler water.

A real distinction between a control system and a natural mechanism such as these, is that:

- A control system can effectively cause the negative entropy, the opposite of natural processes to take place.

That is, a control system can fight against the natural process reaching equilibrium and create a condition that is not otherwise present from physics alone. Of course to do this takes energy, because the overall system is still physical in nature, and so as a whole it still has to meet the laws of thermodynamics.

Natural control systems that fight against entropy

Every cell or other arrangement of matter that we call a life form fights against entropy by creating or enforcing a structure of some sort. A simple example is the structure of a molecule that is a composition of atoms in a structure held together by the 'bonds' between atoms. In essence, three atoms (e.g., one oxygen and two hydrogen) coming together by happenstance in the randomness of space when energy has them moving around, bond together because opposite charges attract and electrons have a 'negative' charge while protons have a 'positive' charge. These hydrogen atoms each contribute their single electron to a shared pool with the oxygen, and the resulting molecule is more stable than the hydrogen or oxygen alone. The lower total energy (entropy) of the molecule keeps it from unbinding without external force, but the molecule is more energetic than the system would be if the atoms never found each other. It sounds a little bit like some marriages. H₂O (two hydrogen and one oxygen as a molecule) is called water. But it's still not a control system. It's just atoms forming a molecule with a locally minimum energy level.

As larger molecules were produced and came together, through randomness at first, they eventually (over many millions of years) formed other structures. Some of these structures ended up being able to reproduce in the proper environment where specific other atoms and molecules were present in adequate quantities and proximity, at the right temperature, pressure, and other conditions. When I say reproduce, I mean that in the proper environment, where there was one such molecule or other structure, components of the composite in proximity ended up resulting in another similar structure which, under similar circumstances, produced similar results. In the capability maturity model this would be the repeatable state of maturity. The net effect of this is that these composite structures are alive, in the sense that they consume energy and other surrounding material components, reproduce, and produce waste in the form of residual components and heat.

Once reproduction starts, the population of these organisms tends to grow based on available supplies (food), so we then get living organisms that turn the general randomness of the environmental components into repetitive patterns of structures, which gets us to the net effect of decreasing entropy locally (in the presence of the increasing number of non-random components in the composite organisms themselves) while of course increasing randomness of the overall system by turning more energy into waste and heat. But the waste itself tends to not be as random as the environment before the living organism was present, because it tends to be specific structures produced as side effects of the reproductive process, which since it is repeated, tends to produce many of the same byproducts. These component byproducts may then, by chance, be composed into other living organisms that themselves reproduce, again altering the environment, and consuming more energy and producing more waste. These organisms that reproduce can be said to:

- **Sense** in that they alter in their chemical and movement of charges (electrical) makeup as they encounter components in the environment.
- **Actuate** in the sense that they ultimately assemble components into composites through their reproductive processes and produce waste components.
- **Communicate** in the sense that the shapes of their surfaces combined with the charges at their bond points attract matching components over a distance (also actuation) and in the sense that they transmit heat changes and waste components into the environment and receive waste components from the environment, some of which may be used for reproduction.
- The question of whether there is control at this point is one that I cannot clearly answer yet, because I don't know enough, and it is pretty close to the edge. The answer lies in whether they respond to senses to alter actuation through communication and whether they respond to their own actuation as it changes the environment and reaches their senses. Nothing about 'intent' is required for control, only perhaps some sort of repeatability so that similar stimulus leads to similar response.

As the environment of multiple living organisms evolved over time, through random variation and selective survival (as an assumption for now), these organisms became more complex, with some mechanisms surviving and reproducing at higher rates, others being consumed as food for others, and so forth. Over the span of time, they eventually developed into cellular structures, presumably as living components combined into larger living composites. These composites came to reproduce differently in several ways.

One of the key parts of cellular reproduction was the development of cells splitting to reproduce. This process involves: collecting, over time, the components required for reproduction, within the cell walls; making a completed reproduction of some internal mechanisms; and then having the cell wall collapse around each of the new smaller cells, making each of them a separate living organism. This is called mitosis, and the different steps in the sequence of reproduction can be understood as a control system because there is a mechanism that changes its interaction (communication) based on sensors to alter the behavior of actuators. As of today, as far as I have been able to figure out, humans do not fully understand how cells work. However:

- **Control** is invoked by the change in phases:⁵
 - In the **growth** phase (called **interphase**), the cell collects components from the environment and allows them to interact with the cell wall to enter, and the DNA structure within the cell reproduces using these components to form a reproduction.
 - Then the cell changes operation into the **Mitotic** phase where the chromosome structures are separated into 2 nuclei:
 - **Prophase**: The chromatin fibers condense into chromosomes that are visible through a light microscope. Each replicated chromosome appears as two identical sister chromatids joined at their centromeres, and the mitotic spindle begins to form. Also, the centrosomes begin to move to opposite poles of the cell, and they are propelled by the lengthening microtubules between them.
 - In **prometaphase**, the nuclear envelope falls apart; microtubules can invade the nuclear area and bind to some of the chromosomes. The microtubules bind at the kinetochores, specialized protein structures at the centromere. Not all microtubules interact with kinetochores. Some microtubules interact with microtubules extending from the other side of the cell.
 - In **metaphase**, the centrosomes have migrated to opposite poles of the cell. The chromosomes have all lined up at the metaphase plate in the middle of the cell, and all chromosomes are attached to microtubules through their kinetochores. The metaphase plate is an imaginary line equidistant from the spindle's 2 poles.
 - In **anaphase**, the shortest stage of mitosis, the sister chromatids break apart, and the chromosomes begin moving to opposite ends of the cell. By the end of anaphase, the 2 halves of the cell have an equivalent collection of chromosomes.
 - In **telophase**, 2 daughter nuclei form. The nuclear envelope begins to reappear. DNA begins to de-condense while spindle microtubules begin to depolymerize. Mitosis, the division of one nucleus into 2, is now complete. Lastly, cytokinesis, which is the division of the cytoplasm, takes place, and the cell divides into 2 separate cells. In animal cells, this is accomplished through a cleavage furrow that pinches the cell in two.

⁵ <https://www.ncbi.nlm.nih.gov/books/NBK482449/> - largely quoted in this explanation with further references to be examined for a more detailed understanding of the control system components and how they work.

- **Regulation** is the name of the internal **control mechanism**. Throughout mitosis, certain checkpoints are essential to the continuation of the process. If certain conditions are not met, mitosis halts. If any of these checkpoints are bypassed without being completed, certain pathologies, such as cancer, can occur:
 - There are three main checkpoints: the G1/S checkpoint, G2/M, and metaphase/anaphase checkpoint.
 - During the G1/S checkpoint, also known as the restriction checkpoint, primary influencers of cell cycle progression include growth factors, DNA damage, cell size, and cell nutrition.
 - The G2/M checkpoint, also known as the DNA replication checkpoint, is influenced by improper DNA replication or DNA damage.
 - The third checkpoint, also known as the spindle apparatus checkpoint, is influenced by the attachment of the mitotic spindle to all chromosomes. Only when all sister chromatids have been bound will mitosis proceed into anaphase.

And thus we have a biological control system operating at the sub-cellular level within the environment of the cell. Arguably, these mechanisms are all just things that emerge from the physics and chemistry of the environment, but the same is true of the toilet. Once the construct exists, it operates because of the forces of nature. Of course toilets don't reproduce.

The hierarchy of operational capabilities

As systems get more complex, there are more complex control issues to consider. Ultimately, the architecture of the system limits what it can and cannot do, but before we can start to look at overall architecture, it's important to understand the parameters and how they work and interact.

As a fundamental, time is always an issue in any control system, because a control communicated to an actuator resulting in changes in the media resulting in a sensor sending communication to the control system takes time, and events in the media that happen in the intervening times between sensation and subsequent actuation may be such that the control mechanism makes 'bad' decisions, creating positive feedback instead of negative feedback, or causing the opposite overall effect the system is designed to or would otherwise produce.

Biological processing and movement rates

Low level movements in the biological realm are far faster than most folks imagine before they start to study the subject.⁶ As a basis for movements, speeds tend to be represented in terms of the length of the organism in the direction of travel, so 1 length per second for a person of length 2 meters is 2 meters per second, or about 2 miles per hour (average walking speed). The fastest runner does 100 meters in just under 9 seconds, or 12 meters per second, about 25 miles per hour (but not for an hour). That's about 6 times body length.

⁶ <https://pdodds.w3.uvm.edu/research/papers/others/2015/meyer-vernet2015a.pdf> Am. J. Phys. 83, 719–722 (2015) "How fast do living organisms move: Maximum speeds from bacteria to elephants and whales", Nicole Meyer-Vernet; Jean-Pierre Rospars.

Various studies have been published, but the fastest reported speed for an organism we have found so far is about 200 bl/s (body lengths per second)⁷ Maximum speed divided by maximum sustainable speed bl/s has been found to be an average of 2-3 for mammals. Fish tend to have maximum speeds of 3-4 times average speed.⁸ Terrestrial and aquatic species have approximately linear maximum speed with body mass with variations over approximately factors of 10 relative to the average linear approximation. This goes from body masses ranging from 10^{-16} kg to 10^6 kg. Obviously flying animals travel at far higher speeds relative to body mass and length, but this is not related directly to the rate of motion, rather to their ability to fly. The motion of their wings can be quite fast compared to human motions, and there are lots of other similar things you might want to look into for better understanding.

Motion in cellular processes is, of course, astonishing from the perspective of large animals like people. For example, synthesis in cellular processes involves, in some cases, rotational machines that spin 134 times per second generating 3 molecules of Adenosine triphosphate (ATP) per rotation.⁹ (ATP is the primary energy currency of cells. It is produced through a process called cellular respiration, which occurs in the mitochondria of eukaryotic cells and the cytoplasm of prokaryotic cells.) That's about half the number of rotations per second of a Boeing 737 jet engine. Rotation for motion in the biological realm is identified as from a bacterial flagellum which is present in about half of bacteria species.¹⁰ Regulation of chromosome speeds in mitosis appears to be a complex control system within cells that sets different travel rates for different chromosomes depending on size and travel distance.¹¹ The recent understanding of how Spirostomum is able to move so quickly¹² shows it can shrink to about 40% of its flat 4mm length in a few milliseconds, an acceleration of up to (200 (m/s)^2) .¹³

To the extent these motions are part of a control system, the control system has to carry out sensing, communications, control mechanisms, communication, and actuation in sequence in order to make the next change in behavior. Of course simple control systems like this are often components in larger control systems.

- Local control systems can keep the motor running at a speed by rapidly effecting low-level changes to keep operations within a range of set points. These local control systems are often components in larger composite control systems that adapt the set points of 'lower level' control systems.
- Most multi-cellular organisms have higher level control systems that control higher levels of the composite, and the resulting behavior emerges from the composite.

Composites of composites layer up to bodies and brains of larger and more complex species with each level of composite examining a smaller set of signals over a longer time frame with focus of attention limiting what they can observe and controllably effect. The speed of each of these components limits the speed of the composite. And the architecture of the composite limits what can happen in parallel or in sequence, while the nature of computation limits the capabilities of different architectures for solving different analytical tasks.

7 <https://worldbuilding.stackexchange.com/questions/59709/how-fast-can-a-microscopic-creature-travel>

8 <https://hal.science/hal-01384319v2/document>

9 <https://thebiologist.rsb.org.uk/biologist-features/the-speed-of-life>

10 https://en.wikipedia.org/wiki/Rotating_locomotion_in_living_systems

11 <https://pmc.ncbi.nlm.nih.gov/articles/PMC4578309/pdf/nihms522013.pdf>

12 <https://www.science.org/doi/epdf/10.1126/sciadv.add6550>

13 <https://news.gatech.edu/news/2018/08/07/worlds-fastest-creature-may-also-be-one-smallest>

Computer processing and movement rates

Computer processing seems very fast at the instruction level, but it typically takes quite a few instructions in sequence before sensory input turns into an act through a control system.

A simple transistor control loop

The basic mechanism underlying most computers today is a transistor or some similar semiconductor. Semiconductors have the property that, by adding impurities (called doping), their electrical characteristics are changed.¹⁴ Doping can increase the number of 'free' protons or electrons, and more free particles allow more free flow of electrical current. At a junction between two different 'dopings', they operate differently in one direction of current flow than the other. In the simplest configuration, a 'diode' can be used to force current to flow only in one direction from one of two 'leads' through the device. This is similar to a semi-permeable membrane found in the biological, chemical, and materials world. Instead of fluids or gasses going one way and not the other, in a diode, current flows in one direction and not the other.¹⁵

A more important configuration for digital systems has three leads, and current coming in one lead (the base) causes increased current flow between the other two leads (the collector and emitter). This acts as an amplifier so that the control current amplifies the controlled current. The speed with which these effects take place is essentially limited by the speed with which the free electrons can move through the semi-conductor substrate. When one electron is put in one side of the device, another comes out the other side, so the time frame for signal propagation has to do with the propagation of the electromagnetic wave through the device and not the movement of individual particles from one side to the other through the device. The 'propagation delay'¹⁶ through a transistor can be as little as a picosecond (10^{-12} seconds a.k.a. ps), which ultimately means that it can process changes in the control signal with this delay between inbound sensor communication and outbound communication to actuators.

Communications at the level of electronics operate at the speed of light in the media, which in wire is about 1ft/nanosecond (10^{-9} sec =1 ns), so an input wire of 1/1000 of a foot, or 1/300th of a meter (3 millimeters: mm) takes about 1 ps to traverse. The fastest such a system could realistically operate would be something like 1 ps for each of (1) actuator (e.g., light emitting diode), (2) transit through media to a sensor (e.g., light sense transistor), (3) turn sensed phenomena into electrical signal, (4) transit to the transistor, (5) the transistor generates the actuator signal change, and (6) transmission over the wire to the actuator. It's a total of about 6 ps for the control loop to control the light level of the LED at a preset level defined by the input to the control transistor. If we reduce the wires and media to 0 length, it's still 3ps

More complex analog control systems can be built using more transistors and resulting longer travel distances between transistors, resulting in longer control cycles. If time is really tight, we could operate a whole set of these small control mechanisms in close proximity, and time offset from each other to control multiple emitters or have multiple feeds to a single emitter, and have feedback from different transistors connected to the same sensor or multiple sensors. But this is a practical limit of how tight an electronic control loop can realistically be. We might be able to scale it down a bit, of course, but the size of atoms required is limiting.

14 [https://en.wikipedia.org/wiki/Doping_\(semiconductor\)](https://en.wikipedia.org/wiki/Doping_(semiconductor))

15 https://en.wikipedia.org/wiki/Semipermeable_membrane

16 https://en.wikipedia.org/wiki/Propagation_delay

A digital switching system

Digital systems typically use more than one transistor to hold a single bit of memory. If we use technologies like metal oxide semiconductor field effect transistors (MOSFETs) we can store a single bit on a single transistor, but even then getting the value in and out of the transistor requires more circuitry. But let's start with the control function, which itself has to take digital inputs and translate them into digital outputs. In digital systems, we call circuits that perform functions that turn a set of input bits into a set of output bits such that the same input bits produce the same output bits every time, 'gates'.

	00	01	11	10
00	1	1	1	1
01	0	0	1	0
11	0	0	1	0
10	1	1	1	1

Gates perform functions like 'AND', 'OR', and 'NOT' at the simplest level, with each combination of inputs producing a desired output. Each function has a delay between input changes and output changes we will call a 'gate delay'. Gate delays in modern digital systems range from picoseconds to more than 10 nanoseconds, and can be faster or slower depending on temperature, technology, power consumption, and other factors. We call the transform from input to output a "transfer function".

We can write the transfer function as a table taking input bits and producing an output bit, with one such transfer function for each output bit. As an example (at left), 4 input bits produce 16 different output combinations. An input of 01 (column) 10 (row) yields an output of 1, while 0101 yields a 0. Assume a 'NOT' gate is inherent in circuit design.

Different circuitry can be used to accomplish this, but as a general principle, there is a time/space tradeoff associated with such systems. Any input to output transfer function can be performed using a single level 'sum of products' circuit that is comprised of an 'OR' gate fed by a set of 'AND' gates, each 'AND' gate taking inputs required for one of the locations in the table. For example, an AND gate taking NOT 1st bit AND 2nd bit AND 3rd bit, and NOT 4th bit (location 0110) in the table yields a '1' output if all 4 conditions are true, and a 0 otherwise. The OR of all of the ANDs (one per table location) will then yield the desired resulting output. The time to execute this function will be 2 gate delays (one AND level and one OR level). A gate delay is a function of the technology used to implement the gates, and for now, the wire lengths will be ignored. In this case, 16 AND gates and one OR gate, each AND gate taking 4 inputs and the OR gate taking 16 inputs produce the output. A more efficient version would do ANDs of inputs [1 and 2], [3 and not 4], and [not 3 and not 4] and OR them together, taking 3 2-input AND gates and one 3-input OR gate. The execution time will still be 2 gate delays.

This is the static case, where inputs stay indefinitely and outputs can be taken way later. The result takes as long as it takes for the inputs to be turned into outputs, and if an input switches, the output might change from one value to another until the inputs propagate through all the gates. So a 1 might flip to a zero, back to a one, and finally settle on a 0, depending on timing. Gate delays are dependent on the number of gates attached to the output of the previous gate, because these delays are usually caused by the amount of time it takes to build up enough charge or current at the inputs to trigger the internal circuitry to cause the internal state change producing the output. The maximum number of connected gates is called 'fan out' (for the output of a gate) and fan-in (for the input) and generally assumes it has to meet the performance specifications connected to gates of the same type.

Control systems with combinational logic only

Using combinational logic alone, a digital control system can take input from a sensor communicated as values over a set of communication lines, and produce a set of output values communicated to sensors over a set of communications lines. By the fact that these are combinational logic circuits only, they have no memory, and cannot handle sequences of bits transferred over a single communications channel. They need as many communications lines as there are input and output bits. Within a computer, these sorts of mechanisms exist, but only inside closely connected gates performing a single function at high speed. High speed, because operating in this mode, all signals arrive in parallel, changes happen as they happen, and outputs go out as soon as they are available. Since timing is important to getting the right answer, there are special circuit types, such as 'self-timed' circuits¹⁷ that are designed to operate cleanly this way, and more generally, these are 'asynchronous circuits'¹⁸ that are not synchronized and must be designed more carefully to avoid so-called 'race conditions'¹⁹ where data values race each other through the circuits.

Storage and sequential finite state machines

Circuitry can also implement storage components that store bits as voltage level or other physical conditions. These stored values then feed inputs to combinational logic elements that have their results stored in other storage components once they are ready. The way most modern computers determine when the outputs are ready is by having a clock that runs at a predetermined rate, slow enough for the combinational circuits to complete their tasks and have stable outputs ready to store. This is usually indicated as the 'clock speed' of the processor. If you go too slow, the computer will operate slowly, and if you go too fast, some of the outputs will not settle down in time to have valid outputs and you lose reliability.

State	Input	Output	Next state
A	1001	1	A
A	0111	0	B
B	1100	0	B
B	1111	1	A

These finite state machines store some number of bits of 'state' and, in a clock cycle, take stored and external inputs, run them through combinational logic, and produce outputs and subsequent (i.e., next) state values that are then stored for the next step in sequence. The example at left shows a description of a finite state machine (FSM) with two states (called A and B), 4 bits of input, one bit of output, and a next state for each entry. Note there are no entries for some inputs in some states, and if this is the specification, the machine might do something unexpected or undesired in the unspecified situations, depending on implementation.

Finite state machines can implement complex algorithms involving multiple steps, producing new output values to the control system only after performing the intermediate calculations. The more steps and the slower the clock cycle, the longer it takes to complete the control loop and send output to actuators in response to input from sensors. Just as communications limit the amount of content that can be sent to a control component, the control process and processor speed limit the speed with which controls can act.

17 <https://ntrs.nasa.gov/api/citations/20040068174/downloads/20040068174.pdf>

18 https://en.wikipedia.org/wiki/Asynchronous_circuit

19 https://en.wikipedia.org/wiki/Race_condition

Computer hardware

Computers are predominantly composed of FSMs. These FSMs take “inputs” (bits) in context of their current “state” (stored as bits) and produce “outputs” (bits) and updated “states”. The “states” are typically stored in memory components similar to the memory components in random access memory, but typically not arranged in the same manner. Inputs are electrically coded signals (they are coded as voltages or similar physical phenomena), typically in one of two acceptable ranges of voltages indicating one of two states per input wire (one bit per wire) or similar transport structure, as are outputs. The transformation from input and current state to output and next state are done by combinational logic circuits that produce different electrically encoded binary values depending on the state and input. These machines are typically operated in synchronization with a “clock” indicating when the operation should be undertaken to store the new values based on old values and when the input and output are stabilized to the point where they can be reliably used. Clock speeds can range up to something like 10 gate delays. For the fastest custom transistor circuits directly attached to each other, perhaps you could get to 30 ps.

Computers combine many of these FSMs to perform higher level functions, many of which operate in parallel and in synchronization with each other to perform the tasks comprising the overall function of the computer or device. Thus most computers operate with many parallel components forming many parallel FSMs. These FSMs communicate both directly to each other and through structures called a “bus”, one or more of which may be present depending on the particulars of the computer architecture. A bus is usually a shared set of “wires” that connect to many internal FSMs allowing one or more to connect to the wires at a time and induce or interpret the sequences of voltages or signals being sent. For example, local storage areas called “registers” containing typically from 8 to 64 bits of binary memory, might be connected to the bus and read or written by any of a number of different FSMs also connected to the bus at different times. The same sorts of structures may be used to connect other FSMs or devices within the computer system and many such busses may be present and simultaneously operable.

If a control process involves complex sequences of FSMs, each settling to a result before the next is invoked, this takes more time than a simpler control system, and again limits the ability of the control mechanism to some resulting pace of transforming changes in input to changes in output. In most industrial control systems, the need for timing constraints on physical elements of the process produces local FSMs in a Programmable Logic Controller (PLC) that implement set points or bounds on input parameters from sensors and produce output to actuators to keep the media within the desired bounds, while using a ‘higher level’ slower process (usually called a Supervisory Control And Data Acquisition (SCADA) system) to determine and implement changes in bounds and set points. Multiple layers of hierarchy can also be used to control higher level functions of the overall system, such as an assembly line controlling machines performing individual steps, controlling individual sets of sensors and actuators, in a Distributed Control System (DCS).²⁰ These systems use Human Computer Interface (HCI) systems to interface to human controllers and retain information about past process events in a history mechanism often called a “Data Historian”. Communications in these systems usually use specific protocols designed for process controls that are usually quite different from those used for other computer communications.

²⁰ https://en.wikipedia.org/wiki/Industrial_control_system

Instruction Sets and General Purpose Computers

While each of these various elements of a control system may be implemented by special purpose devices that only do what they were explicitly designed to do, in almost all cases SCADA, Historian, and HMI interfaces operate on general purpose computers that are capable of doing an unlimited number of different functions. These general purpose computers operate with stored “instructions” formed from an instruction set. The instructions are typically stored in a large array of memory storage devices associated with a random access memory and/or a cache memory, and/or disk storage, and/or network accessible storage. The finite state machines of a central processing unit (CPU) typically:

- “fetch” an instruction from a memory location specified by a “program counter” (PC) register,
- interpret the content of that memory location as an “instruction”,
- sequence a series of actions by different FSMs within the computer to carry out the instruction,
- store any results in various locations in registers, FSMs, and/or memory, and then
- update the PC register to point to the next instruction location.

One instruction after another is processed in this manner to cause a long sequence of complex operations to be undertaken so as to produce more complicated functional behaviors by the computer system. Instructions also sometimes retrieve content from or store content in memory, update the PC register to “jump” to a different location for the next program instruction, or invoke more complex methods that change registers, memory locations, and program execution flow in more complex ways.

Different instructions are typically represented in different ways and may have different storage requirements. For example, an instruction that compares a register to zero and skips the next instruction if they match might be entirely encoded within a single byte (8 bits), while another instruction that adds a value stored anywhere in random access memory (RAM) to a value stored in a register would almost certainly take several bytes just to reference the RAM address.

Because of the sequences of instructions executed by these machines, they tend to operate far more slowly than the FSMs discussed earlier. The fastest single processors today operate at about 9Ghz, which is to say, the clock cycles in just under 10^{-10} seconds, a.k.a. 0.1 ns. Of course these are highly parallelized devices in which each clock cycle doesn’t fetch and execute a full complex instruction, but in which a parallel pipeline is used to divide an instruction into a series of FSMs each feeding the next and executing only a part of each instruction, all operating in parallel. The fastest operations per second (OPS) for a single processor may be in the 10^{10} OPS range, but each instruction itself takes a lot longer from start to finish.

Using multiple processors of this sort operating in parallel (simultaneously) can generate enormous numbers of OPS (usually tracked as Floating Point OPS, or FLOPS), with the fastest modern computers operating in the 3 exaFLOPS (3×10^{18}) range, but for control mechanisms, the time that matters is the time to execute the required sequence of instructions to turn sensor inputs into actuator outputs, and that remains longer than 1 ns.

Larger control systems

Biological systems for fish, fowl, mammals, plants, and other life forms²¹ now existing or potentially evolvable or creatable, execute higher level control mechanisms for movement, like nerves, muscles controlled by reflexes, learned sequences of motions, coordinated acts, and so forth. These sorts of hierarchical control systems have sensors (nerves) that use chemical processes to produce electrical signals that travel (communication) to local (control) mechanisms that produce reflexive responses by sending (communications) electrical signals to muscles (actuators) that constrict to respond.²² “In the spinal mammal it has been shown that this interval is very brief, amounting only to about 0.001 seconds” (1ms).²³ This varies by a number of environmental conditions, including temperature, healing from wounds, and so forth. This overall reflex time is the combination of different times associated with nerve conduction, sensor ending latency, muscular latency, and synapse delay. To get a sense of this, the blink of a human eye takes about 10ms to start, but before the lid is closed, takes from 40-200ms.²⁴ The reflex time itself is far less than the time to complete the movement. Many other reflexes involve more or less complex activities, but this should help to get a sense of what these times are composed of and how they are broken down.

It is also important to differentiate reflexes from responses, in that reflexes are local to the sensor and actuator, while higher level responses to stimuli go through the brain. The reflexes that are local, such as frog leg movements, go through about 3 neurons in the local control system, and those take about 0.7 ms each. Brain cells (neurons) take about 10-50 ms to ‘sum’ incoming signals from dendrites (input structures) and trigger output signals. Typical outputs are 0.5-2 ms bursts of electrical signals, as pulsed sequences.²⁵

Complex thought processes require many neurons and thus take longer and longer time frames to produce more complex results. If this starts to sound similar to the way computers work, welcome to the world of cybernetics. It seems inherent in the nature of things.

Timing and the Implications to physical structures

Time is related to distance through the speed of light. It acts as a speed limit you cannot get around no matter how hard you push the accelerator. That also means that, at the limits of performance, the time to calculate output of the control mechanism from inputs and states is limited by distance between components. An interesting recent example is in financial trading systems. These systems make automated decisions about when to move money from one stock, bond, or other instrument into a cash account, and from a cash account into another fungible instrument. Because the value of these change at every transaction, getting in faster means buying before the price goes up or selling before the price goes down. For that reason, there is a need for speed. The folks making these transactions locate their mechanisms as close as they can to the systems doing the transactions, because the speed of light makes a difference in which trade gets in first or last. The competition is so tight that systems are now using short wave radio instead of fiber optics to get their signals through faster.²⁶

21 [https://en.wikipedia.org/wiki/Taxonomy_\(biology\)](https://en.wikipedia.org/wiki/Taxonomy_(biology))

22 <https://www.sciencedirect.com/topics/biochemistry-genetics-and-molecular-biology/reflex-time>

23 <https://royalsocietypublishing.org/doi/epdf/10.1098/rspb.1921.0004>

24 <https://en.wikipedia.org/wiki/Blinking>

25 <https://www.biosourcesoftware.com/post/the-brain-in-time>

26 <https://www.radioworld.com/tech-and-gear/market-makers-want-to-expand-their-use-of-shortwave>

Operating parameter limitations

The density with which we can pack the components is also limited by their operating parameters, in particular, power and temperature requirements. Each component takes in energy, performs computation, and emits heat and possibly other waste products. As the heat builds up, the components reach their maximum operating temperature and ultimately fail if it gets too hot. There are many schemes for pulling out the heat, but they ultimately involve decreasing the density to make space for the cooling mechanism. At some density, we can pack no more components, and speed must be sacrificed.

One of the key things to understand about the physical architecture of control systems is that there are performance advantages of placement and routing of components and signals. This has long been studied in integrated circuit design and automated routing and placement systems are used in almost all processor designs today.²⁷ Biological systems developed in this way as well, and if you look at reflexes as local control systems, they are faster because of proximity and limited control mechanisms. Responses requiring the central nervous system are further away from limbs but closer to eyes and ears, and they tend to go through more steps and take longer than strictly local reflexes, but they are used for coordinated activities.

Interference limitations

If it weren't for this heat problem, we could pack circuits more closely, but that might also lead to other issues. In particular, electrical systems with current flows and voltages across barriers interact (called coupling) through:

- **Inductance** (the change in charge over time produce effects on other charges at distance that reduces as the square of the distance) so as they get closer, currents in one component create or alter currents in other nearby components; and
- **Capacitance** (charges on one side of a barrier cause attraction of opposite and repulsion of like charges on the other side of the barrier), and as the barrier gets thinner, the charges have more effect across the barrier.

These effects are stronger the closer components are, and at some point, pattern sensitive faults occur, producing wrong results because of interference at small distances. This transfer of electromagnetic energy through space over distances produces all sorts of problems in getting electrical components too close to each other. For example, even if one wire in close proximity changing at a very high rate does not cause enough inductance in another wire close to it, as we put more wires closer, the effects adds up, and if a whole bunch of wires all go from a lower voltage to a higher voltage at the same time, their effects at distance adds up. In digital systems this is particularly problematic because once something gets over the threshold from a 1 to a 0 or a 0 to a 1, the mechanisms use positive feedback to force them further, whereas in analog systems, small changes tend to produce more small changes. Of course wave effects allow small changes to add up in analog systems as well.

Also note that wires of different lengths will better amplify signals at different frequencies. This is how antennas are made, and every part of an electrical circuit acts like a sort of antenna whether we want it to or not. So there are possibilities of interactions from signals at longer distances if the frequencies are right for the physical shapes of the components.

²⁷ [https://en.wikipedia.org/wiki/Placement_\(electronic_design_automation\)](https://en.wikipedia.org/wiki/Placement_(electronic_design_automation))

Multiple frequencies within composites

Yes, it gets more complicated... If multiple signals at different frequencies are around, then they form 'beat frequencies' that can result from wave addition and subtraction. In essence, just like waves in water add up, so do electromagnetic waves in space, and just like physical structures in water towers tend to produce synchronization at various frequencies based on their sizes and shapes, so do physical structures in electronic systems.

- As sizes get smaller and components get closer, less power is required to move fewer charges shorter distances, operations can go faster, and higher frequency interactions become more complex and difficult to manage, but ...
- As sizes get bigger and components get further away, more power is required to move more charges longer distances, lower frequency interactions are more important and complex, and longer wavelength antennae are in place picking up different signals, often from further away.

Composites almost always have components that operate in different time frames, producing more frequencies of signals and a more complex electromagnetic environment.

Timing and coordination at distance

Composites such as those described here involve more complex coordination as they get more complex and encounter more complex external conditions. For example the timing of activities producing a spinning mechanism at the cellular level has to be right or the spinning element will not spin and produce the compounds it is supposed to produce, and the metabolism of the entire cellular organism will ultimately fail.

This is just like the timing of spark plugs in a gasoline engine. If the sparks come at the right time, they explode the gasoline so it expands as the cylinders are moving to expand the chamber volume, and it adds energy to the rotation of the drive shaft. But if you explode as the cylinders are reducing the available volume, you are likely to have the engine seize up and perhaps do more extensive damage (blow up).

This is just like the coordination required to do any other complex tasks involving multiple actuators working together. If the actions are the same but with different timing, the system will not operate properly. Timing is controlled by the feedback system operating in concert with the physicality of the sensors, actuators, communication, and control components in terms of their time characteristics. Learning systems, whether in computers, biological systems, or otherwise, learn timing or sequencing as part of their process for coordination. This typically happens by trying different timings in sequence, and finding out which ones work, which ones don't, and adjusting the control parameters to perform sequences of steps with timing. They learn to anticipate where a target of an act will be and aim ahead of it.

A typical system of this sort will develop sequences of internal steps that track sensors awaiting conditions consistent with the conditions that were successful before and initiating next actions to the actuators at times that worked previously. Additional experimentation by small increments helps adapt the process to changing conditions and seeks to optimize the timing for those specific sequences. Babies of many species learn to walk and adapt their walking as their limbs grow longer, learn to leap, stand, and do so with different external conditions, all by coordinating timing at a distance.

Shielding

Designers and biological systems face complexity levels so high that no amount of available computing power or analysis or controls can accurately model effects of design or operational decisions on electromagnetic interference, thermodynamics, or other dynamic interactions. Fortunately, designers know about shielding and nature has evolved it. We may think we are more clever than nature, but shielding is present in biological systems as well.

Just as some materials or configurations of materials can be used to carry or amplify electrical, magnetic, sonic, chemical, or thermal effects, other materials or configurations of materials tend to capture or block these effects. Glass is pretty good at holding most chemical compounds, and as a practical matter it is used in some chemical processes for allowing optical information to pass while restricting chemical flows. Cell walls are pretty good at keeping molecular interference down within cells, and they usually allow only the right stuff in.

In the same way, metal enclosures tend to reduce or eliminate electromagnetic transmissions through the air. Unless they are at frequencies that end up treating the metal as antennae or end up physically touching.

- Insulators, such as ceramics, are used for high voltage lines to separate them from the poles they are connected to, preventing short circuits.
- Insulation on wires tends to work against physical connections but tends to be ineffective against electromagnetic waveform emanations, so we see foil or other very thin shielding around the individual wire insulators to prevent waveform leakage, and another layer of insulator outside the metal foil to protect against the shielding acting as a conductor or being damaged. Shielding of this sort works both ways, so it reduces interference from external signals on the internal wires as well as leakage from the wires being emitted. The internal wires are also twisted to reduce the crosstalk between different pairs enclosed in the same shielding, yielding the shielded twisted pair wiring technology in widespread use for electronics such as telephone wiring.²⁸
- Heat shielding is used to reduce the transfer of heat between different components. A great example is vacuum separation used in thermos bottles. But heat shielding can only slow the transfer, and not eliminate it. Unless there is some way for the heat to escape, it will just keep building up. So heat has to be removed as well as shielded.
- Light shielding prevents the passage of optical signals, generally photons, from passing a barrier, and is used, for example, to isolate fiber optic signals from the external environment. However, because most fiber optic connections use laser light, dispersal is less of an issue than interference of other sorts. Sun glasses and other similar shielding is used for people to prevent their eyes from becoming blinded by the light, for example in welding or nuclear tests.

Biological systems also have various forms of shielding. This ranges from cell walls that shield the internal works of cells from external molecules, to sheaths on nerves and muscles that prevent interference in layers of protection, to semi-permeable membranes within cells and larger creatures, to skin and other separators that coat different organs. Eye lids, nails, and ultimately, even clothing and nests act as shields against intrusions of different sorts of compounds and intrusions.

²⁸ https://en.wikipedia.org/wiki/Twisted_pair

Central vs Distributed systems

Centralized control systems have the property that all of the signals from all of the sensors, all of the signals to all of the actuators, and all of the stored information are available for control in one place and all the time. They also have the property that all of those sensor signals get there in a time related to whatever distance the central control is from each of them, that the outbound signals to actuators get there in a time related to whatever distance they each are from the control system, and that all of those signals have to be processed by the centralized control mechanism. If there is a failure in the central control mechanism, it has the potential to effect all of the operations of the entire system. It has all of the complexity of all of the control mechanisms for all of the components of the composite.

Distributed control systems have the property that overall control is shared between more than one control mechanism. That allows, for example, local control over 'reflex' responses at higher speed if there is physical proximity, preprocessing of local signals with less bandwidth required for controls further away from the sensors and actuators, less local processing power in each control mechanism, operation of the local mechanisms when other local mechanisms fail, less heat, less interference, and fewer problems associated with larger centralized control systems. However, signals and information are not then all available at a single place for coordination and decision-making.

Parallel processing

As identified earlier, many FSMs operate simultaneously within computer systems and even within individual components of those systems. In this sense, parallel processing has long been understood and applied to resolve performance problems, to allow and support expandability of systems, and as a core part of computer system architecture. For example, every Ethernet component added to a computer system operates in parallel with other Ethernet components in that system, each typically with its own processor for processing Ethernet packets. At every level of the hardware within computer systems, this is done, from the internal components that comprise the "adders" and "multipliers" used to do simple internal calculations, to the memory elements within computer systems such as registers and the internal states of each FSM, to the multitude of I/O devices found in common computer mother boards, to the disk drives, serial ports, multiple CPUs, and so forth. In that sense, parallelism has existed and been exploited in design since the earliest days of computers.

There is nothing new about the use of parallelism and many parallel processing structures, architectures, components, computers, and configurations have long existed. But of course in the biological realm, every cell operates in parallel with other cells; in the chemical realm, every chemical reaction happens in parallel with other chemical reactions; and every atom operates in parallel with every other atom. All of the activity that ever happens, happens at the same time as other activities, and the control mechanisms of all of the composites of these components are present at multiple levels of structure.

In networked environments, there is inherent parallelism because each computer or other device connected to networks, including the devices used to transport network traffic, are essentially all operating in parallel all of the time. If this were not true, when I turned off my local computer, the rest of the Internet would stop. Of course as far as I can tell, the Internet is not working when I cannot observe it or the side effects of its operations. Don't blink!

Communications and control architecture

The properties of communications architecture have, over time, led to overall structures:

- **At the most minute levels of physics and chemistry**, control appears to be non-existent or extremely local. For example, while all of the statistical effects of all of the electrons in the universe theoretically effect each other electron with a time delay of the speed of light, in combination, all of these effects are negligible at the level of how electrons remain in orbit around nuclei and how they operate as part of chemical bonds forming molecules.
- **At the cellular level**, chemical processes largely operate in parallel, however control mechanisms that have been identified are local to the cell, effected by surrounding regions of space and their properties such as temperature, pressure, acidity, etc..
- **At the transistor level in electronics**, control mechanisms are strictly local to the transistor, with one lead controlling flow between two other leads in the simple cases most commonly encountered, but environmental conditions effect the mechanisms and within the transistor, atomic processes act in parallel.
- **At the gate level in electronics**, everything tends to be asynchronous and control mechanisms are implemented locally by the transistors or other active components that operate the gate and tend to settle in digital states instead of acting as analog mechanisms. Internally parallel, but externally sequential.
- **At the reflex level in biological systems** controls tend to be distributed in nature, with each reflex operating independently, even if overridable by control signals (inhibitory axons)²⁹ from higher levels of the composite organism they are part of.
- **At the FSM level in electronics**, except when using self-timed circuits, operation tends to be synchronized by clocks, and for that reason, there is clock centralization, while each FSM component in the composite operates in parallel with other FSMs producing a distributed system composite.
- **At the response level in biological systems**, neural pathways tend to reach a central control mechanism in a star configuration with all local control systems connecting to the overall central control (i.e., central nervous) system of the composite 'creature'.
- **At the computer level in electronics**, there are usually central processing units controlling local input and output devices and memory, these days with multiple 'cores' acting as a localized parallel processing facility, so the overall structure is a local network of processors controlling a set of peripheral devices in a 2-level hierarchy.
- **At the level of the 'creature'**, there tend to be hierarchical systems of controls that will be discussed in later chapters in more depth, but where each hierarchical element is a parallel processing composite of networked components, each with internal control mechanisms and overarching control mechanisms that effect each other through things like chemical concentrations. Everything is parallel, self-timed or asynchronous.
- **At the network level in electronics**, systems tend to be distributed hierarchies, but this is quite complicated and non-uniform, and will be discussed in later chapters.

²⁹ <https://pmc.ncbi.nlm.nih.gov/articles/PMC2885981/>

Up the cognitive ladder

As control system composites become larger aggregates of more components, an apparently emergent effect is an increasing level of cognizance (defined roughly in dictionaries as knowledge, awareness, or notice) in the overall composite. Note the link between the concept of focus of attention and 'notice'. Without commenting on things like awareness, knowledge, or self-awareness, system structures built up by nature or design tend to form more complex and highly interacting structures in order to control larger more complex composites in more complex environments over longer time frames.

At a simplistic level, biological systems tend to have more complex 'brains' as they control more complex bodies in more complex situations, and automated control systems tend to have more complex 'computers' as they control more complex mechanisms in more complex environments.

There tends to be a hierarchical control mechanism at some level of locality, and a networked mechanism at a higher level of complexity, and a hierarchical control system above that and a networked control system above that, and so forth. Hierarchy has the advantage of centralization for efficient coordination and localization of failures, while networked distributed control has the advantage of speed and reduced complexity for things that can be simplified to that level of operation, and overall resilience.

At different levels of complexity, networked systems break down and are unable to survive, leading to hierarchies that coordinate more effectively. But the centralization leads to lack of performance required for survival, leading to a higher level of distributed systems that again run into complexity problems, and so forth.

This seems to operate all the way up to the social structures of animals in groups, and for now, up to human and computer integrated global society.

Fundamental control system limits

This set of limitations of control systems appears to be an inherent result of the nature of time and space in control systems. While we can find a wide range of ways to architect improvements in one area or another, at the end of the day, physics limits densities by energy and heat constraints, time scales by the speed of light and chemical and electrical processes in the components, and cognitive capacities by the focus of attention limitations, and the compression of physical realities into the detectable and processable information and content.

This has implications for making such systems workable, reliable, and resilient, but also in terms of the ways these systems can fail.