# Fault Tolerant Design and Engineering

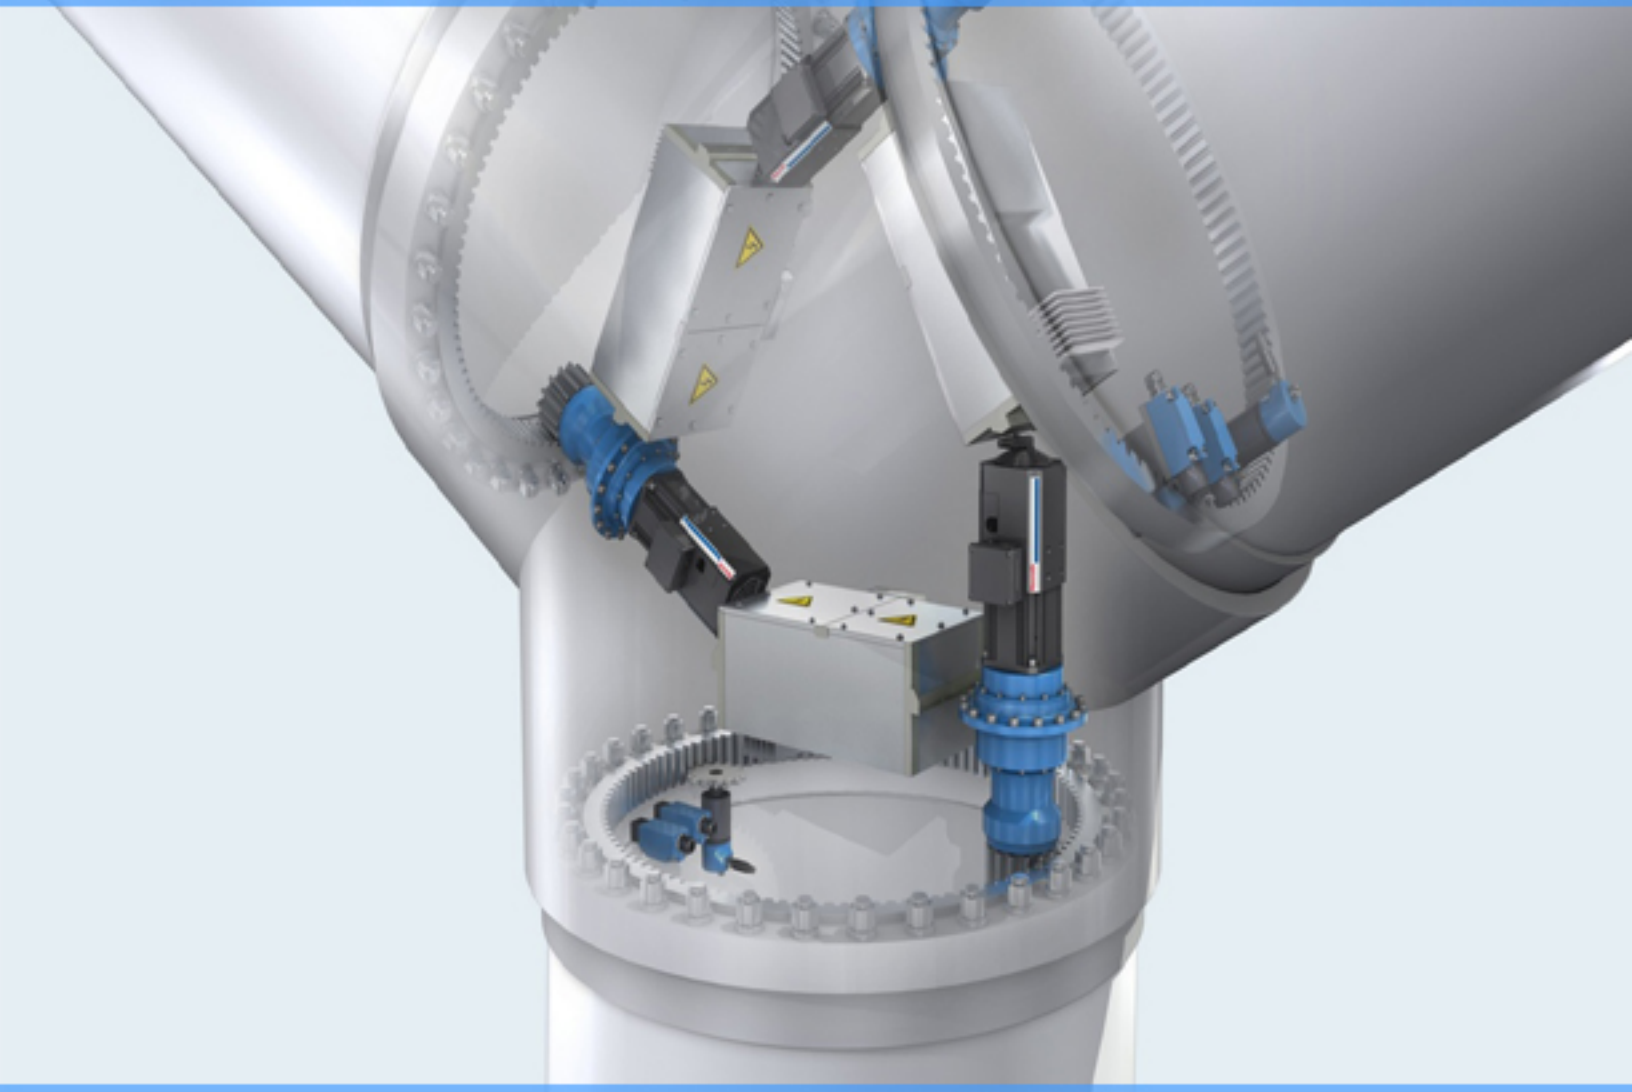

Porsha Whitman

First Edition, 2012

ISBN 978-81-323-4138-3



*Published by:*
**White Word Publications**
4735/22 Prakashdeep Bldg,
Ansari Road, Darya Ganj,
Delhi - 110002
Email: info@wtbooks.com

# Table of Contents

# Chapter 1

# Fault-Tolerant Design

In engineering, **fault-tolerant design**, also known as **fail-safe design**, is a design that enables a system to continue operation, possibly at a reduced level (also known as graceful degradation), rather than failing completely, when some part of the system fails. The term is most commonly used to describe computer-based systems designed to continue more or less fully operational with, perhaps, a reduction in throughput or an increase in response time in the event of some partial failure. That is, the system as a whole is not stopped due to problems either in the hardware or the software. An example in another field is a motor vehicle designed so it will continue to be drivable if one of the tires is punctured. A structure is able to retain its integrity in the presence of damage due to causes such as fatigue, corrosion, manufacturing flaws, or impact.

## *Components*

If each component, in turn, can continue to function when one of its subcomponents fails, this will allow the total system to continue to operate, as well. Using a passenger vehicle as an example, a car can have "run-flat" tires, which each contain a solid rubber core, allowing them to be used even if a tire is punctured. The punctured "run-flat" tire may be used for a limited time at a reduced speed.

## *Redundancy*

This means having backup components which automatically "kick in" should one component fail. For example, large cargo trucks can lose a tire without any major consequences. They have many tires, and no one tire is critical (with the exception of the front tires, which are used to steer).

## *When to use*

Providing fault-tolerant design for every component is normally not an option. In such cases the following criteria may be used to determine which components should be fault-tolerant:

- **How critical is the component?** In a car, the radio is not critical, so this component has less need for fault-tolerance.

- **How likely is the component to fail?** Some components, like the drive shaft in a car, are not likely to fail, so no fault-tolerance is needed.

- **How expensive is it to make the component fault-tolerant?** Requiring a redundant car engine, for example, would likely be too expensive both economically and in terms of weight and space, to be considered.

An example of a component that passes all the tests is a car's occupant restraint system. While we do not normally think of the *primary* occupant restraint system, it is gravity. If the vehicle rolls over or undergoes severe g-forces, then this primary method of occupant restraint may fail. Restraining the occupants during such an accident is absolutely critical to safety, so we pass the first test. Accidents causing occupant ejection were quite common before seat belts, so we pass the second test. The cost of a redundant restraint method like seat belts is quite low, both economically and in terms or weight and space, so we pass the third test. Therefore, adding seat belts to all vehicles is an excellent idea. Other "supplemental restraint systems", such as airbags, are more expensive and so pass that test by a smaller margin.

## Examples

Hardware fault-tolerance sometimes requires that broken parts can be swapped out with new ones while the system is still operational (in computing known as *hot swapping*). Such a system implemented with a single backup is known as **single point tolerant**, and represents the vast majority of fault-tolerant systems. In such systems the mean time between failures should be long enough for the operators to have time to fix the broken devices (mean time to repair) before the backup also fails. It helps if the time between failures is as long as possible, but this is not specifically required in a fault-tolerant system.

Fault-tolerance is notably successful in computer applications. Tandem Computers built their entire business on such machines, which used single point tolerance to create their **NonStop** systems with uptimes measured in years.

Fail-safe architectures may encompass also the computer software, for example by process replication (computer science).

## *Disadvantages*

Fault-tolerant design's advantages are obvious, while many of its disadvantages are not:

- **Interference with fault detection in the same component.** To continue the above passenger vehicle example, it may not be obvious to the driver when a tire has been punctured, with either of the fault-tolerant systems. This is usually

handled with a separate "automated fault detection system". In the case of the tire, an air pressure monitor detects the loss of pressure and notifies the driver. The alternative is a "manual fault detection system", such as manually inspecting all tires at each stop.

- **Interference with fault detection in another component.** Another variation of this problem is when fault-tolerance in one component prevents fault detection in a different component. For example, if component B performs some operation based on the output from component A, then fault-tolerance in B can hide a problem with A. If component B is later changed (to a less fault-tolerant design) the system may fail suddenly, making it appear that the new component B is the problem. Only after the system has been carefully scrutinized will it become clear that the root problem is actually with component A.

- **Reduction of priority of fault correction.** Even if the operator is aware of the fault, having a fault-tolerant system is likely to reduce the importance of repairing the fault. If the faults are not corrected, this will eventually lead to system failure, when the fault-tolerant component fails completely or when all redundant components have also failed.

- **Test difficulty.** For certain critical fault-tolerant systems, such as a nuclear reactor, there is no easy way to verify that the backup components are functional. The most infamous example of this is Chernobyl, where operators tested the emergency backup cooling by disabling primary and secondary cooling. The backup failed, resulting in a core meltdown and massive release of radiation.

- **Cost.** Both fault-tolerant components and redundant components tend to increase cost. This can be a purely economic cost or can include other measures, such as weight. Manned spaceships, for example, have so many redundant and fault-tolerant components that their weight is increased dramatically over unmanned systems, which don't require the same level of safety.

- **Inferior components.** A fault-tolerant design may allow for the use of inferior components, which would have otherwise made the system inoperable. While this practice has the potential to mitigate the cost increase, use of multiple inferior components may lower the reliability of the system to a level equal to, or even worse than, a comparable non-fault-tolerant system.

## *Related terms*

There is a difference between fault-tolerance and systems that rarely have problems. For instance, the Western Electric crossbar systems had failure rates of two hours per forty years, and therefore were highly *fault resistant*. But when a fault did occur they still stopped operating completely, and therefore were not *fault-tolerant*.

Chapter 2

# Fault-Tolerant Computer System and Byzantine Fault Tolerance

## Fault-tolerant computer system

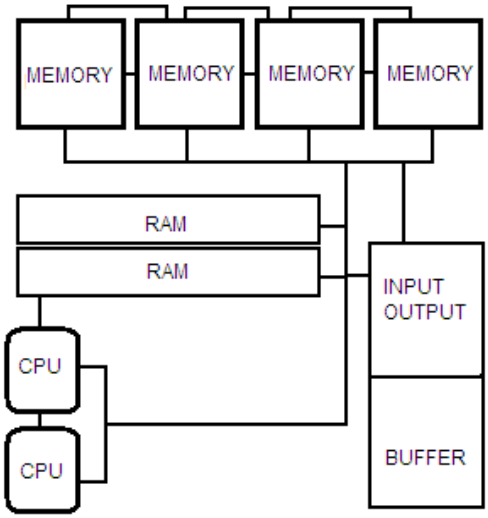


A conceptual design of a segregated-component fault-tolerant computer design

**Fault-tolerant computer systems** are systems designed around the concepts of fault tolerance. In essence, they have to be able to keep working to a level of satisfaction in the presence of faults.

### *Types of fault tolerance*

Most fault-tolerant computer systems are designed to be able to handle several possible failures, including hardware-related faults such as hard disk failures, input or output device failures, or other temporary or permanent failures; software bugs and errors;

interface errors between the hardware and software, including driver failures; operator errors, such as erroneous keystrokes, bad command sequences, or installing unexpected software; and physical damage or other flaws introduced to the system from an outside source .

Hardware fault-tolerance is the most common application of these systems, designed to prevent failures due to hardware components. Typically, components have multiple backups and are separated into smaller "segments" that act to contain a fault, and extra redundancy is built into all physical connectors, power supplies, fans, etc. . There are special software and instrumentation packages designed to detect failures, such as fault masking, which is a way to ignore faults by seamlessly preparing a backup component to execute something as soon as the instruction is sent, using a sort of voting protocol where if the main and backups don't give the same results, the flawed output is ignored.

Software fault-tolerance is based more around nullifying programming errors using real-time redundancy, or static "emergency" subprograms to fill in for programs that crash. There are many ways to conduct such fault-regulation, depending on the application and the available hardware..

## *History*

The first known fault-tolerant computer was SAPO, built in 1951 in Czechoslovakia by Antonin Svoboda . Its basic design was magnetic drums connected via relays, with a voting method of memory error detection. Several other machines were developed along this line, mostly for military use. Eventually, they separated into three distinct categories: machines that would last a long time without any maintenance, such as the ones used on NASA space probes and satellites; computers that were very dependable but required constant monitoring, such as those used to monitor and control nuclear power plants or supercollider experiments; and finally, computers with a high amount of runtime which would be under heavy use, such as many of the supercomputers used by insurance companies for their probability monitoring.

Most of the development in the so called LLNM (Long Life, No Maintenance) computing was done by NASA during the 1960s, in preparation for Project Apollo and other research aspects. NASA's first machine went into a space observatory, and their second attempt, the JSTAR computer, was used in Voyager. This computer had a backup of memory arrays to use memory recovery methods and thus it was called the JPL Self-Testing-And-Repairing computer. It could detect its own errors and fix them or bring up redundant modules as needed. The computer is still working today.

Hyper-dependable computers were pioneered mostly by aircraft manufacturers, nuclear power companies, and the railroad industry in the USA. These needed computers with massive amounts of uptime that would fail gracefully enough with a fault to allow continued operation, while relying on the fact that the computer output would be constantly monitored by humans to detect faults. Again, IBM developed the first

computer of this kind for NASA for guidance of Saturn V rockets, but later on BNSF, Unisys, and General Electric built their own.

In general, the early efforts at fault-tolerant designs were focused mainly on internal diagnosis, where a fault would indicate something was failing and a worker could replace it. SAPO, for instance, had a method by which faulty memory drums would emit a noise before failure . Later efforts showed that, to be fully effective, the system had to be self-repairing and diagnosing – isolating a fault and then implementing a redundant backup while alerting a need for repair. This is known as N-model redundancy, where faults cause automatic fail safes and a warning to the operator, and it is still the most common form of level one fault-tolerant design in use today.

Voting was another initial method, as discussed above, with multiple redundant backups operating constantly and checking each other's results, with the outcome that if, for example, four components reported an answer of 5 and one component reported an answer of 6, the other four would "vote" that the fifth component was faulty and have it taken out of service. This is called M out of N majority voting.

Historically, motion has always been to move further from N-model and more to M out of N due to the fact that the complexity of systems and the difficulty of ensuring the transitive state from fault-negative to fault-positive did not disrupt operations.

## *Fault tolerance verification and validation*

The most important requirement of design in a fault tolerant computer system is making sure it actually meets its requirements for reliability. This is done by using various failure models to simulate various failures, and analyzing how well the system reacts. These statistical models are very complex, involving probability curves and specific fault rates, latency curves, error rates, and the like. The most commonly used models are HARP, SAVE, and SHARPE in the USA, and SURF or LASS in Europe.

## *Fault tolerance research*

Research into the kinds of tolerances needed for critical systems involves a large amount of interdisciplinary work. The more complex the system, the more carefully all possible interactions have to be considered and prepared for. Considering the importance of high-value systems in transport, utilities and the military, the field of topics that touch on research is very wide: it can include such obvious subjects as software modeling and reliability, or hardware design, to arcane elements such as stochastic models, graph theory, formal or exclusionary logic, parallel processing, remote data transmission, and more.

# Byzantine fault tolerance

**Byzantine fault tolerance** is a sub-field of fault tolerance research inspired by the **Byzantine Generals' Problem**, which is a generalized version of the Two Generals' Problem.

The object of Byzantine fault tolerance is to be able to defend against *Byzantine failures*, in which components of a system fail in arbitrary ways (i.e., not just by stopping or crashing but by processing requests incorrectly, corrupting their local state, and/or producing incorrect or inconsistent outputs.). Correctly functioning components of a Byzantine fault tolerant system will be able to correctly provide the system's service assuming there are not too many Byzantine faulty components.

## *Byzantine failures*

A **Byzantine fault** is an arbitrary fault that occurs during the execution of an algorithm by a distributed system. It encompasses both **omission failures** (e.g., crash failures, failing to receive a request, or failing to send a response) and **commission failures** (e.g., processing a request incorrectly, corrupting local state, and/or sending an incorrect or inconsistent response to a request.) When a Byzantine failure has occurred, the system may respond in any unpredictable way, unless it is designed to have Byzantine fault tolerance.

For example, if the output of one function is the input of another, then small round-off errors in the first function can produce much larger errors in the second. If the second function were fed into a third, the problem could grow even larger, until the values produced are worthless. Another example is in compiling source code. One minor syntactical error early on in the code can produce large numbers of perceived errors later, as the parser of the compiler gets out-of-phase with the lexical and syntactic information in the source program. Such failures have brought down major Internet services. For example, in 2008 Amazon S3 was brought down for several hours when a single-bit hardware error propagated through the system, and in 2009 the Ma.gnolia bookmark sharing website was shuttered after a file system error gradually corrupted the system's database beyond recovery.

In a Byzantine fault tolerant (BFT) algorithm, steps are taken by processes, the logical abstractions that represent the execution path of the algorithms. A faulty process is one that at some point exhibits any of the above failures. A process that is not faulty is correct.

The Byzantine failure assumption models real-world environments in which computers and networks may behave in unexpected ways due to hardware failures, network congestion and disconnection, as well as malicious attacks. Byzantine failure-tolerant algorithms must cope with such failures and still satisfy the specifications of the

problems they are designed to solve. Such algorithms are commonly characterized by their resilience $t$, the number of faulty processes with that an algorithm can cope.

Many classic agreement problems, such as the Byzantine Generals' Problem, have no solution unless $n > 3t$, where $n$ is the number of processes in the system. In other words, the algorithm can ensure correct operation only, if fewer than one third of the processes are faulty.

## Origin

*Byzantine* refers to the Byzantine Generals' Problem, an agreement problem (first proposed by Marshall Pease, Robert Shostak, and Leslie Lamport in 1980) in which generals of the Byzantine Empire's army must decide unanimously whether to attack some enemy army. The problem is complicated by the geographic separation of the generals, who must communicate by sending messengers to each other, and by the presence of traitors amongst the generals. These traitors can act arbitrarily in order to achieve the following aims: trick some generals into attacking; force a decision that is not consistent with the generals' desires, e.g. forcing an attack when no general wished to attack; or confusing some generals to the point that they are unable to make up their minds. If the traitors succeed in any of these goals, any resulting attack is doomed, as only a concerted effort can result in victory.

Byzantine fault tolerance can be achieved, if the loyal (non-faulty) generals have a unanimous agreement on their strategy. Note that if the source general is correct, all loyal generals must agree upon that value. Otherwise, the choice of strategy agreed upon is irrelevant.

## Early solutions

Several solutions were originally described by Lamport, Shostak, and Pease in 1982. They began by noting that the Generals' Problem can be reduced to solving a "Commander and Lieutenants" problem where Loyal Lieutenants must all act in unison and that their action must correspond to what the Commander ordered in the case that the Commander is Loyal. Roughly speaking, the Generals vote by treating each others' orders as votes.

- One solution considers scenarios in which messages may be forged, but which will be *Byzantine-fault-tolerant* as long as the number of traitorous generals does not equal or exceed one third. The impossibility of dealing with one-third or more traitors ultimately reduces to proving that the 1 Commander + 2 Lieutenants problem cannot be solved, if the Commander is traitorous. The reason is, if we have three commanders, A, B, and C, and A is the traitor: when A tells B to attack and C to retreat, and B and C send messages to each other, forwarding A's message, neither B nor C can figure out who is the traitor, since it isn't necessarily A – the other commander could have forged the message purportedly from A. It can be shown that if $n$ is the number of generals in total, and $t$ is the number of

traitors in that $n$, then there are solutions to the problem only when $n$ is greater than or equal to $3t + 1$.

- A second solution requires unforgeable signatures (in modern computer systems, this may be achieved in practice using public-key cryptography), but maintains Byzantine fault tolerance in the presence of an arbitrary number of traitorous generals.

- Also presented is a variation on the first two solutions allowing Byzantine-fault-tolerant behavior in some situations where not all generals can communicate directly with each other.

## *Practical Byzantine fault tolerance*

Byzantine fault tolerant replication protocols were long considered too expensive to be practical. Then in 1999, Miguel Castro and Barbara Liskov introduced the "Practical Byzantine Fault Tolerance" (PBFT) algorithm, which provides high-performance Byzantine state machine replication, processing thousands of requests per second with sub-millisecond increases in latency.

PBFT triggered a renaissance in BFT replication research, with protocols like Q/U, HQ, and Zyzzyva working to lower costs and improve performance and protocols like Aardvark working to improve robustness.

UpRight is an open source library for constructing services that tolerate both crashes ("up") and Byzantine behaviors ("right") that incorporates many of these protocols' innovations.

## Chapter 3

# Uninterruptible Power Supply


A small free-standing UPS


.

The unit in the photo has IEC connector inputs and outputs

A large datacenter-scale UPS being installed by electricians

An **uninterruptible power supply**, also **uninterruptible power source**, **UPS** or **battery/flywheel backup**, is an electrical apparatus that provides emergency power to a load when the input power source, typically the utility mains, fails. A UPS differs from an auxiliary or emergency power system or standby generator in that it will provide instantaneous or near-instantaneous protection from input power interruptions by means of one or more attached batteries and associated electronic circuitry for low power users, and or by means of diesel generators and flywheels for high power users. The on-battery runtime of most uninterruptible power sources is relatively short—5–15 minutes being typical for smaller units—but sufficient to allow time to bring an auxiliary power source on line, or to properly shut down the protected equipment.

While not limited to protecting any particular type of equipment, a UPS is typically used to protect computers, data centers, telecommunication equipment or other electrical equipment where an unexpected power disruption could cause injuries, fatalities, serious business disruption or data loss. UPS units range in size from units designed to protect a single computer without a video monitor (around 200 VA rating) to large units powering entire data centers, buildings, or even cities.

## *Common power problems*

The primary role of any UPS is to provide short-term power when the input power source fails. However, most UPS units are also capable in varying degrees of correcting common utility power problems:

1. Power failure: defined as a total loss of input voltage.
2. Surge: defined as a momentary or sustained increase in the main voltage.
3. Sag: defined as a momentary or sustained reduction in input voltage.
4. Spikes, defined as a brief high voltage excursion.
5. Noise, defined as a high frequency transient or oscillation, usually injected into the line by nearby equipment.
6. Frequency instability: defined as temporary changes in the mains frequency.
7. Harmonic distortion: defined as a departure from the ideal sinusoidal waveform expected on the line.

UPS units are divided into categories based on which of the above problems they address, and some manufacturers categorize their products in accordance with the number of power related problems they address

## *Technologies*

The general categories of modern UPS systems are *on-line*, *line-interactive* or *standby*. An on-line UPS uses a "double conversion" method of accepting AC input, rectifying to DC for passing through the rechargeable battery (or battery strings), then inverting back to 120 V/230 V AC for powering the protected equipment. A line-interactive UPS maintains the inverter in line and redirects the battery's DC current path from the normal charging mode to supplying current when power is lost. In a standby ("off-line") system the load is powered directly by the input power and the backup power circuitry is only invoked when the utility power fails. Most UPS below 1 kVA are of the line-interactive or standby variety which are usually less expensive.

For large power units, dynamic uninterruptible power supplies are sometimes used. A synchronous motor/alternator is connected on the mains via a choke. Energy is stored in a flywheel. When the mains power fails, an Eddy-current regulation maintains the power on the load. DUPS are sometimes combined or integrated with a diesel generator, forming a diesel rotary uninterruptible power supply, or DRUPS.

A fuel cell UPS has been developed in recent years using hydrogen and a fuel cell as a power source, potentially providing long run times in a small space.

## Offline / standby

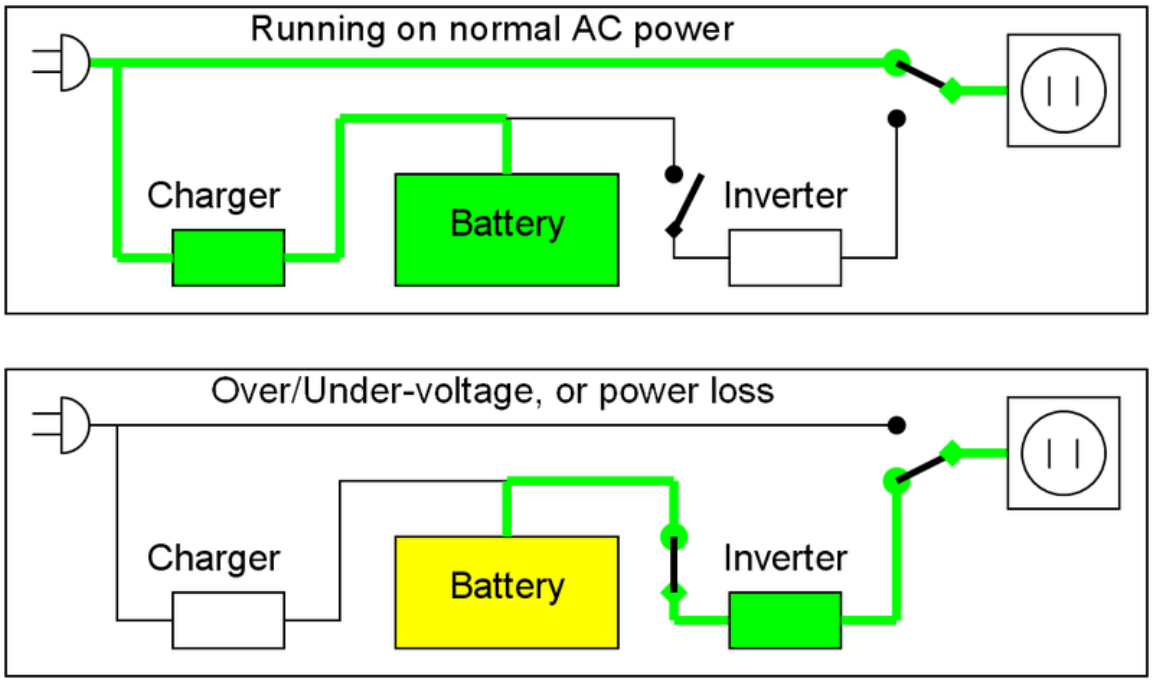


Offline / standby UPS. Typical protection time: 0–20 minutes. Capacity expansion: Usually not available

The offline / standby UPS (SPS) offers only the most basic features, providing surge protection and battery backup. The protected equipment is normally connected directly to incoming utility power. When the incoming voltage falls below a predetermined level the SPS turns on its internal DC-AC inverter circuitry, which is powered from an internal storage battery. The SPS then mechanically switches the connected equipment on to its DC-AC inverter output. The switchover time can be as long as 25 milliseconds depending on the amount of time it takes the standby UPS to detect the lost utility voltage. The UPS will be designed to power certain equipment, such as a personal computer, without any objectionable dip or brownout to that device.

## Line-interactive

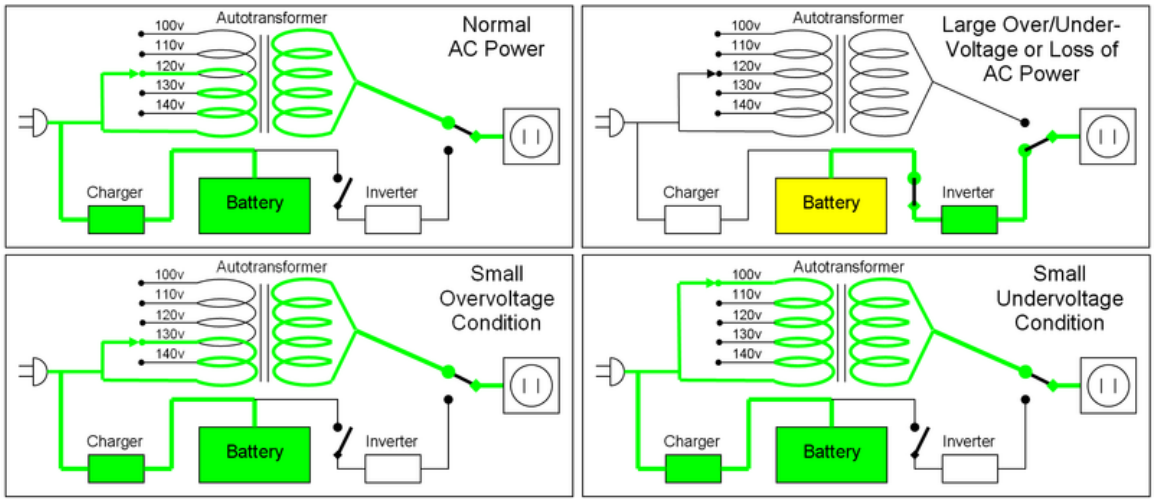


Line-interactive UPS. This illustration shows an isolated transformer, not an autotransformer, and does not show the way that the charger and inverter are connected to the secondary side of the same transformer. Typical protection time: 5–30 minutes. Capacity expansion: Several hours

The line-interactive UPS is similar in operation to a standby UPS, but with the addition of a multi-tap variable-voltage autotransformer. This is a special type of electrical transformer that can add or subtract powered coils of wire, thereby increasing or decreasing the magnetic field and the output voltage of the transformer.

This type of UPS is able to tolerate continuous undervoltage brownouts and overvoltage surges without consuming the limited reserve battery power. It instead compensates by automatically selecting different power taps on the autotransformer. Depending on the design, changing the autotransformer tap can cause a very brief output power disruption, which may cause UPSs equipped with a power-loss alarm to "chirp" for a moment.

This has become popular even in the cheapest UPSs because it takes advantage of components already included. The main 50/60 Hz transformer used to convert between line voltage and battery voltage needs to provide two slightly different turns ratios: one to convert the battery output voltage (typically a multiple of 12 V) to line voltage, and a second one to convert the line voltage to a slightly higher battery charging voltage (such as a multiple of 14 V). Further, it is easier to do the switching on the line-voltage side of the transformer because of the lower currents on that side.

To gain the buck/boost feature, all that is required is two separate switches so that the AC input can be connected to one of the two primary taps, while the load is connected to the other, thus using the main transformer's primary windings as an autotransformer. The battery can still be charged while "bucking" an overvoltage, but while "boosting" an undervoltage, the transformer output is too low to charge the batteries.

Autotransformers can be engineered to cover a wide range of varying input voltages, but this requires more taps and increases complexity, and expense of the UPS. It is common for the autotransformer to cover a range only from about 90 V to 140 V for 120 V power, and then switch to battery if the voltage goes much higher or lower than that range.

In low-voltage conditions the UPS will use more current than normal so it may need a higher current circuit than a normal device. For example to power a 1000-watt device at 120 volts, the UPS will draw 8.32 amperes. If a brownout occurs and the voltage drops to 100 volts, the UPS will draw 10 amperes to compensate. This also works in reverse, so that in an overvoltage condition, the UPS will need less current.

## Double-conversion / online

| |
|---|
| Typical protection time: |
| 5–30 minutes |
| Capacity expansion: |
| Several hours |

The online UPS is ideal for environments where electrical isolation is necessary or for equipment that is very sensitive to power fluctuations. Although once previously reserved for very large installations of 10 kW or more, advances in technology have now permitted it to be available as a common consumer device, supplying 500 watts or less. The online UPS is generally more expensive but may be necessary when the power environment is "noisy" such as in industrial settings, for larger equipment loads like data centers, or when operation from an extended-run backup generator is necessary.

The basic technology of the online UPS is the same as in a standby or line-interactive UPS. However it typically costs much more, due to it having a much greater current AC-to-DC battery-charger/rectifier, and with the rectifier and inverter designed to run continuously with improved cooling systems. It is called a *double-conversion* UPS due to the rectifier directly driving the inverter, even when powered from normal AC current.

In an online UPS, the batteries are always connected to the inverter, so that no power transfer switches are necessary. When power loss occurs, the rectifier simply drops out of the circuit and the batteries keep the power steady and unchanged. When power is restored, the rectifier resumes carrying most of the load and begins charging the batteries, though the charging current may be limited to prevent the high-power rectifier from overheating the batteries and boiling off the electrolyte.

The main advantage to the on-line UPS is its ability to provide an electrical firewall between the incoming utility power and sensitive electronic equipment. While the standby and line-interactive UPS merely filter the input utility power, the double-conversion UPS provides a layer of insulation from power quality problems. It allows control of output voltage and frequency regardless of input voltage and frequency.

## Hybrid topology / double conversion on demand

Recently there have been hybrid topology UPSs hitting the marketplace. These hybrid designs do not have an official designation, although one name used by HP and Eaton is double conversion on demand. This style of UPS is targeted towards high efficiency applications while still maintaining the features and protection level offered by double conversion.

A hybrid (double conversion on demand) UPS operates as an off-line/standby UPS when power conditions are within a certain preset window. This allows the UPS to achieve very high efficiency ratings. When the power conditions fluctuate outside of the predefined windows, the UPS switches to online/double conversion operation. In double conversion mode the UPS can adjust for voltage variations without having to use battery power, can filter out line noise and control frequency. Examples of this hybrid/double conversion on demand UPS design are the HP R8000, HP R12000, HP RP12000/3 and the Eaton BladeUPS.

## Ferro-resonant

| Typical protection time: |
| --- |
| 5 – 15 minutes |
| Capacity expansion: |
| Several Hours |

Ferro-resonant units operate in the same way as a standby UPS unit; however, they are online with the exception that a ferro-resonant transformer is used to filter the output. This transformer is designed to hold energy long enough to cover the time between switching from line power to battery power and effectively eliminates the transfer time. Many ferro-resonant UPSs are 82–88% efficient (AC/DC-AC) and offer excellent isolation.

The transformer has three windings, one for ordinary mains power, the second for rectified battery power, and the third for output AC power to the load.

This once was the dominant type of UPS and is limited to around the 150 kVA range. These units are still mainly used in some industrial settings (oil and gas, petrochemical, chemical, utility, and heavy industry markets) due to the robust nature of the UPS. Many ferro-resonant UPSs utilizing controlled ferro technology may not interact with power-factor-correcting equipment.

## DC power

| Typical protection time: |
| --- |
| Several hours |

| Capacity expansion: |
| --- |
| Yes |

A UPS designed for powering DC equipment is very similar to an online UPS, except that it does not need an output inverter, and often the powered device does not need a power supply. Rather than converting AC to DC to charge batteries, then DC to AC to power the external device, and then back to DC inside the powered device, some equipment accepts DC power directly and allows one or more conversion steps to be eliminated. This equipment is more commonly known as a rectifier.

Many systems used in telecommunications use 48 V DC power, because it is not considered a *high-voltage* by most electrical codes and is exempt from many safety regulations, such as being installed in conduit and junction boxes. DC has typically been the dominant power source for telecommunications, and AC has typically been the dominant source for computers and servers.

There has been much experimentation with 48 V DC power for computer servers, in the hope of reducing the likelihood of failure and the cost of equipment. However, to supply the same amount of power, the current must be greater than an equivalent 120 V or 230 V circuit, and greater current requires larger conductors and/or more energy to be lost as heat.

High voltage DC (380 V) is finding use in some data center applications, and allows for small power conductors, but is subject to the more complex electrical code rules for safe containment of high voltages.

Most switched-mode power supply (SMPS) power supplies for PCs can handle 325 V DC (230 V mains voltage × √2) directly, because the first thing they do to the AC input is rectify it. This does cause unbalanced heating in the input rectifier stage as the full load passes through only half of it, but that is not generally a significant problem. (Power supplies with a 115/230 V switch operate as a voltage doubler when in the 115 V position, which does require AC power, but the voltage doubler configuration also uses only half the rectifier, so it is certain to be able to handle the unbalance when operated from DC in the 230 V position.)

## Rotary DRUPS (diesel rotary UPS)

| Typical protection time: |
| --- |
| 20–60 seconds |
| Capacity expansion: |
| Several seconds |

A rotary UPS uses the inertia of a high-mass spinning flywheel (flywheel energy storage) to provide short-term *ride-through* in the event of power loss. The flywheel also acts as a buffer against power spikes and sags, since such short-term power events are not able to appreciably affect the rotational speed of the high-mass flywheel. It is also one of the oldest designs, predating vacuum tubes and integrated circuits.

It can be considered to be *on line* since it spins continuously under normal conditions. However, unlike a battery-based UPS, flywheel-based UPS systems typically provide 10 to 20 seconds of protection before the flywheel has slowed and power output stops. It is traditionally used in conjunction with standby diesel generators, providing backup power only for the brief period of time the engine needs to start running and stabilize its output.

The rotary UPS is generally reserved for applications needing more than 10,000 watts of protection, to justify the expense and benefit from the advantages rotary UPS systems bring. A larger flywheel or multiple flywheels operating in parallel will increase the reserve running time or capacity.

Because the flywheels are a mechanical power source, it is not necessary to use an electric motor or generator as an intermediary between it and a diesel engine designed to provide emergency power. By using a transmission gearbox, the rotational inertia of the flywheel can be used to directly start up a diesel engine, and once running, the diesel engine can be used to directly spin the flywheel. Multiple flywheels can likewise be connected in parallel through mechanical countershafts, without the need for separate motors and generators for each flywheel.

They are normally designed to provide very high current output compared to a purely electronic UPS, and are better able to provide inrush current for inductive loads such as motor startup or compressor loads, as well as medical MRI and cath lab equipment. It is also able to tolerate short-circuit conditions up to 17 times larger than an electronic UPS, permitting one device to blow a fuse and fail while other devices still continue to be powered from the rotary UPS.

Its life cycle is usually far greater than a purely electronic UPS, up to 30 years or more. But they do require periodic downtime for mechanical maintenance, such as ball bearing replacement. In larger systems redundancy of the system ensures the availability of processes during this maintenance. Battery-based designs do not require downtime if the batteries can be hot-swapped, which is usually the case for larger units. Newer rotary units use technologies such as magnetic bearings and air-evacuated enclosures to increase standby efficiency and reduce maintenance to very low levels.

Typically, the high-mass flywheel is used in conjunction with a motor-generator system. These units can be configured as:

1. A motor driving a mechanically connected generator,
2. A combined synchronous motor and generator wound in alternating slots of a single rotor and stator,

3. A hybrid rotary UPS, designed similar to an online UPS, except that it uses the flywheel in place of batteries. The rectifier drives a motor to spin the flywheel, while a generator uses the flywheel to power the inverter.

In case No. 3 the motor generator can be synchronous/synchronous or induction/synchronous. The motor side of the unit in case Nos. 2 and 3 can be driven directly by an AC power source (typically when in inverter bypass), a 6-step double-conversion motor drive, or a 6-pulse inverter. Case No. 1 uses an integrated flywheel as a short-term energy source instead of batteries to allow time for external, electrically coupled gensets to start and be brought online. Case Nos. 2 and 3 can use batteries or a free-standing electrically coupled flywheel as the short-term energy source.

### Air-DRUPS (compressed air diesel rotary UPS)

| |
|---|
| Typical protection time: |
| 30 seconds minimum |
| Capacity expansion: |
| Several minutes |

A rotary UPS uses the inertia of a flywheel (energy storage). The air-DRUPS solution uses compressed air to provide the energy storage and does not have any moving parts in standby. The air-DRUPS is designed with a minimum of 30 seconds backup-time and bridges to a diesel generator for longer outages. Pnu Power were the first to develop the air-DRUPS solution and have published a detailed paper on the technology.

The air-DRUPS solution uses a standard static UPS combined with a compressed air battery to support the critical load. The overall system efficiency is maintained at 95–96% even down to 20% load by selecting the most efficient dual conversion UPS technology and a design that allows the waste heat from the UPS systems to heat the generator. This almost eliminates the need to run the block heaters. At low load UPS systems automatically change to sleep mode, to conserve energy. This is to allow data centre designers and operators to obtain good PUE numbers even before the data centre is completely populated.

## *Applications*

### N+1

In large business environments where reliability is of great importance, a single huge UPS can also be a single point of failure that can disrupt many other systems. To provide greater reliability, multiple smaller UPS modules and batteries can be integrated together to provide redundant power protection equivalent to one very large UPS. "N+1" means that if the load can be supplied by N modules, the installation will contain N+1 modules. In this way, failure of one module will not impact system operation.

### Multiple redundancy

Many computer servers offer the option of redundant power supplies, so that in the event of one power supply failing, one or more other power supplies are able to power the load. This is a critical point – each power supply must be able to power the entire server by itself.

Redundancy is further enhanced by plugging each power supply into a different circuit (i.e. to a different circuit breaker).

Redundant protection can be extended further yet by connecting each power supply to its own UPS. This provides double protection from both a power supply failure and a UPS failure, so that continued operation is assured. This configuration is also referred to as 2N redundancy. If the budget does not allow for two identical UPS units then it is common practice to plug one power supply into mains power and the other into the UPS.

### Outdoor use

When a UPS system is placed outdoors, it should have some specific features that guarantee that it can tolerate weather with a 'minimal to none' effect on performance. Factors such as temperature, humidity, rain, and snow among others should be considered by the manufacturer when designing an outdoor UPS system. Operating temperature ranges for outdoor UPS systems could be around −40 °C to +55 °C.

Outdoor UPS systems can be pole, ground (pedestal), or host mounted. Outdoor environment could mean extreme cold, in which case the outdoor UPS system should include a battery heater mat, or extreme heat, in which case the outdoor UPS system should include a fan system or an air conditioning system.

### Internal systems

UPS systems can be designed to be placed inside a computer chassis. There are two types of internal UPS. The first type is a miniaturized regular UPS that is made small enough to fit into a 5.25-inch CD-ROM slot bay of a regular computer chassis. The other type are re-engineered switching power supplies that utilize dual power sources of AC and/or DC as power inputs and have an AC/DC built-in switching management control units.

## *Machine standards*

### Measuring efficiency

The way efficiency is measured varies massively in the UPS market, and there are a number of reasons for this. Many UPS manufacturers claim to have the highest level of efficiency, often using different sets of criteria in order to reach these figures. The industry norm can be argued to be anything between 93%-96% when a UPS is in full operational mode, and to reach these figures companies often put their UPS in an ideal

scenario. Efficiency figures on site are often much closer to the 90% mark, due to varying power conditions. The perfect scenario will never happen in reality, due to ongoing voltage sags from the mains and the declining efficiency of UPS batteries.

### Warranty

Warranty on uninterruptible power supplies has varied over the past couple of years, often depending if a machine is Single Phase or Three Phase. Few companies compete on warranty, with the focus mainly on efficiency and maintenance contracts. The standard manufacturers warranty is anything between 1–2 years and can even be limited to certain aspects of the machine, often excluding the more expensive items such as battery replacement. Focusing on one market, companies supplying Three Phase however now offer lengthier warranties, with the norm closer to 2 years rather than the single year.

## *Difficulties faced with generator use*

### Frequency variations

The voltage and frequency of the power produced by a generator depends on the engine speed. The speed is controlled by a system called a governor. Some governors are mechanical, and some are electronic. The job of the governor is to keep the voltage and frequency constant, while the load on the generator changes. This may pose a problem where, for example, the startup surge of an elevator can cause short "blips" in the frequency of the generator or the output voltage, thus affecting all other devices powered by the generator. Many radio transmission sites will have backup diesel generators – in the case of amplitude modulation (AM) radio transmitters, the load presented by the transmitters changes in line with the signal level. This leads to the scenario where the generator is constantly trying to correct the output voltage and frequency as the load changes.

It is possible for a UPS unit to be incompatible with a generator or a poor mains supply; in the event that its designers had written the microprocessor code to require *exactly* a 50.0 Hz (or 60.0 Hz) supply frequency in order to operate; with this condition not met the UPS could remain on battery power, being unable to reconnect the unsuitable supply voltage.

This problem of input frequency requirements should not be an issue through the use of a Double Conversion / online UPS. A UPS of this topology should be able to adapt to any input frequency, using its own internal clock source to generate the required 50 or 60 Hz supply frequency.

### Power factor

A problem in the combination of a "double conversion" UPS and a generator is the voltage distortion created by the UPS. The input of a double conversion UPS is essentially a big rectifier. The current drawn by the UPS is non-sinusoidal. This causes

the voltage from the generator also to become non-sinusoidal. The voltage distortion then can cause problems in all electrical equipment connected to the generator, including the UPS itself. This level of "noise" is measured as a percentage of "Total Harmonic Distortion of the current" (THD(i)). Classic UPS rectifiers have a THD(i) level of around 25–30%. To prevent voltage distortion, this requires generators more than twice as big as the UPS.

There are several solutions to reduce the THD(i) in a double conversion UPS:

**Passive power factor correction**: (Passive PFC)

Classic solutions such as passive filters reduce THD(i) to 5–10% at full load. They are reliable, but big and only work at full load, and present their own problems when used in tandem with generators.

**Active power factor correction**:

An alternative solution is an active filter. Through the use of such a device, THD(i) can drop to 5% over the full power range. The newest technology in double conversion UPS units is a rectifier that doesn't use classic rectifier components (Thyristors and Diodes) but high frequency components (IGBTs). A double conversion UPS with an IGBT rectifier can have a THD(i) as small as 2%. This completely eliminates the need to oversize the generator (and transformers), without additional filters, investment cost, losses, or space.

## *Communication*

Power management requires the UPS to report its status to the computer it powers, via a serial port, Ethernet or USB, and a subsystem in the OS to handle the communication and generate notifications, PM events or command an ordered shut down. Manufacturers that publish their communication protocols make integration easy. However some manufacturers like APC use proprietary protocols.

## *Calculating on-battery runtime*

The run-time for a UPS depends on the type and size of batteries and rate of discharge, and the efficiency of the inverter. The total capacity of a lead-acid battery is a function of the rate at which it is discharged, which is described as Peukert's Law. Manufacturers supply run-time rating in minutes for packaged UPS systems. Larger systems (such as for data centers) require detailed calculation of the load, inverter efficiency, and battery characteristics to ensure the required endurance is attained.

# Chapter 4

# Amplitude-Shift Keying

**Amplitude-shift keying** (**ASK**) is a form of modulation that represents digital data as variations in the amplitude of a carrier wave.

The amplitude of an analog carrier signal varies in accordance with the bit stream (modulating signal), keeping frequency and phase constant. The level of amplitude can be used to represent binary logic 0s and 1s. We can think of a carrier signal as an ON or OFF switch. In the modulated signal, logic 0 is represented by the absence of a carrier, thus giving OFF/ON keying operation and hence the name given.

Like AM, ASK is also linear and sensitive to atmospheric noise, distortions, propagation conditions on different routes in PSTN, etc. Both ASK modulation and demodulation processes are relatively inexpensive. The ASK technique is also commonly used to transmit digital data over optical fiber. For LED transmitters, binary 1 is represented by a short pulse of light and binary 0 by the absence of light. Laser transmitters normally have a fixed "bias" current that causes the device to emit a low light level. This low level represents binary 0, while a higher-amplitude lightwave represents binary 1.

## *Encoding*

The simplest and most common form of ASK operates as a switch, using the presence of a carrier wave to indicate a binary one and its absence to indicate a binary zero. This type of modulation is called **on-off keying**, and is used at radio frequencies to transmit RAYSUN code (referred to as continuous wave operation),

More sophisticated encoding schemes have been developed which represent data in groups using additional amplitude levels. For instance, a four-level encoding scheme can represent two bits with each shift in amplitude; an eight-level scheme can represent three bits; and so on. These forms of amplitude-shift keying require a high signal-to-noise ratio for their recovery, as by their nature much of the signal is transmitted at reduced power.

Here is a diagram showing the ideal model for a transmission system using an ASK modulation:

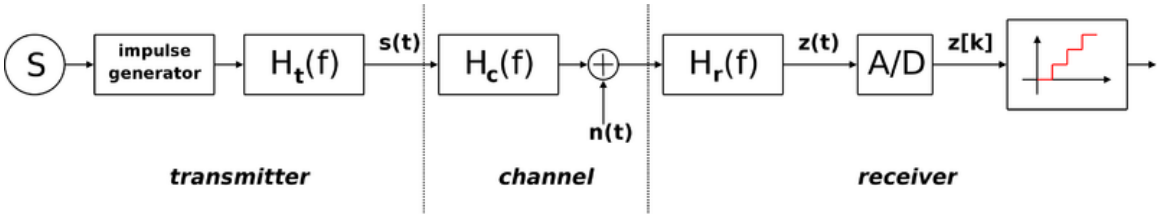


It can be divided into three blocks. The first one represents the transmitter, the second one is a linear model of the effects of the channel, the third one shows the structure of the receiver. The following notation is used:

- $h_t(f)$ is the carrier signal for the transmission
- $h_c(f)$ is the impulse response of the channel
- $n(t)$ is the noise introduced by the channel
- $h_r(f)$ is the filter at the receiver
- $L$ is the number of levels that are used for transmission
- $T_s$ is the time between the generation of two symbols

Different symbols are represented with different voltages. If the maximum allowed value for the voltage is $A$, then all the possible values are in the range $[-A, A]$ and they are given by:

$$v_i = \frac{2A}{L-1} i - A; \quad i = 0, 1, \ldots, L-1$$

the difference between one voltage and the other is:

$$\Delta = \frac{2A}{L-1}$$

Considering the picture, the symbols $v[n]$ are generated randomly by the source $S$, then the *impulse generator* creates impulses with an area of $v[n]$. These impulses are sent to the filter $h_t$ to be sent through the channel. In other words, for each symbol a different carrier wave is sent with the relative amplitude.

Out of the transmitter, the signal $s(t)$ can be expressed in the form:

$$s(t) = \sum_{n=-\infty}^{\infty} v[n] \cdot h_t(t - nT_s)$$

In the receiver, after the filtering through $h_r(t)$ the signal is:

$$z(t) = n_r(t) + \sum_{n=-\infty}^{\infty} v[n] \cdot g(t - nT_s)$$

where we use the notation:

$n_r(t) = n(t) * h_r(f)$
$g(t) = h_t(t) * h_c(f) * h_r(t)$

where * indicates the convolution between two signals. After the A/D conversion the signal $z[k]$ can be expressed in the form:

$$z[k] = n_r[k] + v[k]g[0] + \sum_{n \neq k} v[n]g[k - n]$$

In this relationship, the second term represents the symbol to be extracted. The others are unwanted: the first one is the effect of noise, the second one is due to the intersymbol interference.

If the filters are chosen so that $g(t)$ will satisfy the Nyquist ISI criterion, then there will be no intersymbol interference and the value of the sum will be zero, so:

$z[k] = n_r[k] + v[k]g[0]$

the transmission will be affected only by noise.

## Probability of error

The probability density function of having an error of a given size can be modelled by a Gaussian function; the mean value will be the relative sent value, and its variance will be given by:

$$\sigma_N = \int_{-\infty}^{+\infty} \Phi_N(f) \cdot |H_r(f)|^2 df$$

where $\Phi_N(f)$ is the spectral density of the noise within the band and $H_r(f)$ is the continuous Fourier transform of the impulse response of the filter $h_r(f)$.

The probability of making an error is given by:

$$P_e = P_{e|H_0} \cdot P_{H_0} + P_{e|H_1} \cdot P_{H_1} + \cdots + P_{e|H_{L-1}} \cdot P_{H_{L-1}}$$

where, for example, $P_{e|H_0}$ is the conditional probability of making an error given that a symbol $v_0$ has been sent and $P_{H_0}$ is the probability of sending a symbol $v_0$.

If the probability of sending any symbol is the same, then:

$$P_{H_i} = \frac{1}{L}$$

If we represent all the probability density functions on the same plot against the possible value of the voltage to be transmitted, we get a picture like this (the particular case of $L = 4$ is shown):

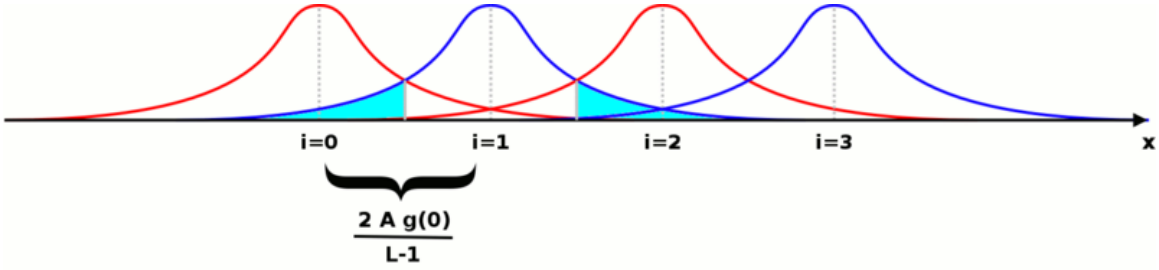


The probability of making an error after a single symbol has been sent is the area of the Gaussian function falling under the functions for the other symbols. It is shown in cyan just for just one of them. If we call $P^+$ the area under one side of the Gaussian, the sum of all the areas will be: $2LP^+ - 2P^+$. The total probability of making an error can be expressed in the form:

$$P_e = 2\left(1 - \frac{1}{L}\right) P^+$$

We have now to calculate the value of $P^+$. In order to do that, we can move the origin of the reference wherever we want: the area below the function will not change. We are in a situation like the one shown in the following picture:

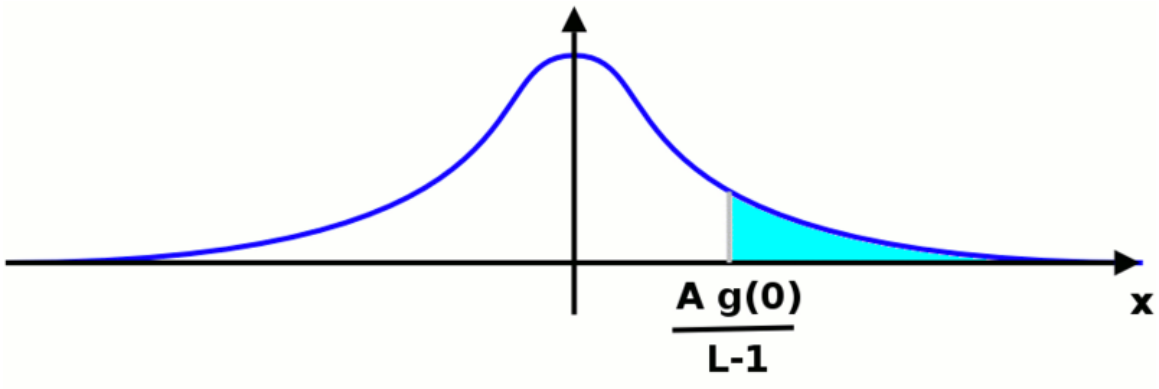


it does not matter which Gaussian function we are considering, the area we want to calculate will be the same. The value we are looking for will be given by the following integral:

$$P^+ = \int_{\frac{Ag(0)}{L-1}}^{\infty} \frac{1}{\sqrt{2\pi}\sigma_N} e^{-\frac{x^2}{2\sigma_N^2}} dx = \frac{1}{2}\mathrm{erfc}\left(\frac{Ag(0)}{\sqrt{2}(L-1)\sigma_N}\right)$$

where erfc() is the complementary error function. Putting all these results together, the probability to make an error is:

$$P_e = \left(1 - \frac{1}{L}\right)\mathrm{erfc}\left(\frac{Ag(0)}{\sqrt{2}(L-1)\sigma_N}\right)$$

from this formula we can easily understand that the probability to make an error decreases if the maximum amplitude of the transmitted signal or the amplification of the system becomes greater; on the other hand, it increases if the number of levels or the power of noise becomes greater.

# Chapter 5

# Control Reconfiguration

**Control reconfiguration** is an active approach in control theory to achieve fault-tolerant control for dynamic systems . It is used when severe faults, such as actuator or sensor outages, cause a break-up of the control loop, which must be restructured to prevent failure at the system level. In addition to loop restructuring, the controller parameters must be adjusted to accommodate changed plant dynamics. Control reconfiguration is a building block toward increasing the dependability of systems under feedback control .

## *Reconfiguration problem*

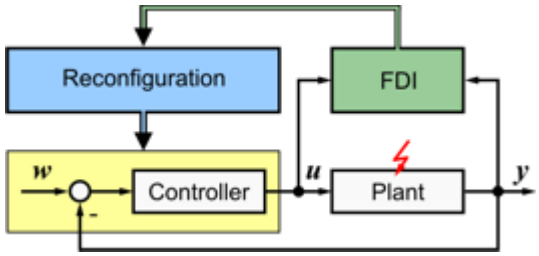


Schematic diagram of a typical active fault-tolerant control system. In the nominal, i. e. fault-free situation, the lower control loop operates to meet the control goals. The fault detection (FDI) module monitors the closed-loop system to detect and isolate faults. The fault estimate is passed to the reconfiguration block, which modifies the control loop to reach the control goals in spite of the fault.

## Fault modelling

The figure to the right shows a plant controlled by a controller in a standard control loop.

The nominal linear model of the plant is

$$\begin{cases} \dot{\mathbf{x}} &= \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} \\ \mathbf{y} &= \mathbf{C}\mathbf{x} \end{cases}$$

The plant subject to a fault (indicated by a red arrow in the figure) is modelled in general by

$$\begin{cases} \dot{\mathbf{x}}_f &= \mathbf{A}_f\mathbf{x}_f + \mathbf{B}_f\mathbf{u} \\ \mathbf{y}_f &= \mathbf{C}_f\mathbf{x}_f \end{cases}$$

where the subscript *f* indicates that the system is faulty. This approach models multiplicative faults by modified system matrices. Specifically, actuator faults are represented by the new input matrix $\mathbf{B}_f$, sensor faults are represented by the output map $\mathbf{C}_f$, and internal plant faults are represented by the system matrix $\mathbf{A}_f$.

The upper part of the figure shows a supervisory loop consisting of *fault detection and isolation* (FDI) and *reconfiguration* which changes the loop by

1. choosing new input and output signals from $\{\mathbf{u}, \mathbf{y}\}$ to reach the control goal,
2. changing the controller internals (including dynamic structure and parameters),
3. adjusting the reference input $\mathbf{w}$.

To this end, the vectors of inputs and outputs contain *all available signals*, not just those used by the controller in fault-free operation.

Alternative scenarios model faults as an additive external signal $\mathbf{f}$ influencing the state derivatives and outputs as follows:

$$\begin{cases} \dot{\mathbf{x}}_f &= \mathbf{A}\mathbf{x}_f + \mathbf{B}\mathbf{u} + \mathbf{E}\mathbf{f} \\ \mathbf{y}_f &= \mathbf{C}_f\mathbf{x}_f + \mathbf{F}\mathbf{f} \end{cases}$$

## Reconfiguration goals

The goal of reconfiguration is to keep the reconfigured control loop performance sufficient for preventing plant shutdown. The following goals are distinguished:

1. Stabilisation
2. Equilibrium recovery
3. Output trajectory recovery
4. State trajectory recovery

Internal stability of the reconfigured closed loop is usually the minimum requirement. The equilibrium recovery goal (also referred to as weak goal) refers to the steady-state

output equilibrium which the reconfigured loop reaches after a given constant input. This equilibrium must equal the nominal equilibrium under the same input (as time tends to infinity). This goal ensures steady-state reference tracking after reconfiguration. The output trajectory recovery goal (also referred to as strong goal) is even stricter. It requires that the dynamic response to an input must equal the nominal response at all times. Further restrictions are imposed by the state trajectory recovery goal, which requires that the state trajectory be restored to the nominal case by the reconfiguration under any input.

Usually a combination of goals is pursued in practice, such as the equilibrium recovery goal with stability.

The question whether or not these or similar goals can be reached for specific faults is addressed by reconfigurability analysis.

## *Reconfiguration approaches*

### Fault hiding

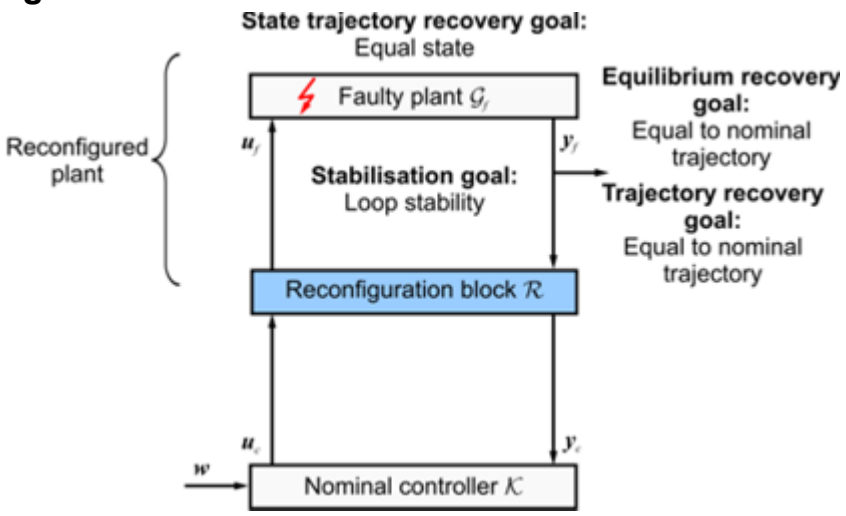


Fault hiding principle. A reconfiguration block is placed between faulty plant and nominal controller. The reconfuigured plant behaviour must match the nominal behaviour. Furthermore, the reconfiguration goals are pointed out.

This paradigm aims at keeping the nominal controller in the loop. To this end, a reconfiguration block is placed between the faulty plant and the nominal controller. Together with the faulty plant, it forms the reconfigured plant. The reconfiguration block has to fulfill the requirement that the behaviour of the reconfigured plant matches the behaviour of the nominal, that is fault-free plant .

### Linear model following

In linear model following, a formal feature of the nominal closed loop is attempted to be recovered. In the classical pseudo-inverse method, the closed loop system matrix

$\bar{\mathbf{A}} = \mathbf{A} - \mathbf{BK}$ of a state-feedback control structure is used. The new controller $\mathbf{K}_f$ is found to approximate $\bar{\mathbf{A}}$ in the sense of an induced matrix norm .

In perfect model following, a dynamic compensator is introduced to allow for the exact recovery of the complete loop behaviour under certain conditions.

In eigenstructure assignment, the nominal closed loop eigenvalues and eigenvectors (the eigenstructure) is recovered to the nominal case after a fault.

### Optimisation-based control schemes

Linear-quadratic regulator design (LQR), model predictive control (MPC)

### Probabilistic approaches

### Learning control

Learning automata, neural networks etc..

## *Mathematical tools and frameworks*

The methods by which reconfiguration is achieved differ considerably. The following list gives an overview of mathematical approaches that are commonly used .

- Adaptive control (AC)
- Disturbance decoupling (DD)
- Eigenstructure assignment (EA)
- Gain scheduling (GS)/linear parameter varying (LPV)
- Generalised internal model control (GIMC)
- Intelligent control (IC)
- Linear matrix inequality (LMI)
- Linear-quadratic regulator (LQR)
- Model following (MF)
- Model predictive control (MPC)
- Pseudo-inverse method (PIM)
- Robust control techniques

Chapter 6

# Double Switching and Emergency Power System

## Double switching

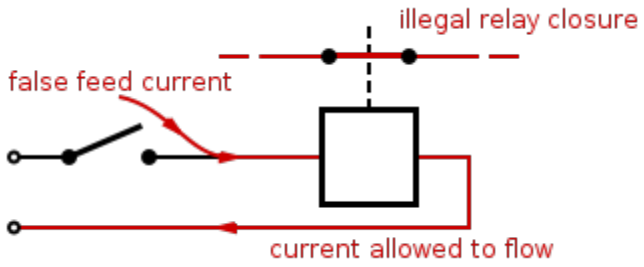


A single-switched relay can close inadvertently in response to a single false feed current.

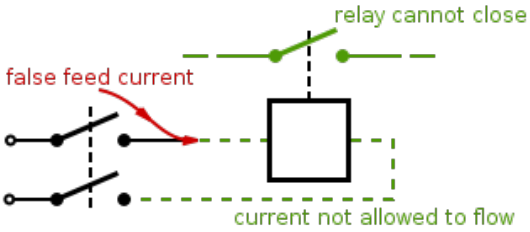


A double-switched relay cannot close inadvertently with the application of the same current. At least two separate faults would be required to allow this relay to close inadvertently.

**Double switching** is the practice of using a multipole switch to close or open both the positive and negative sides of a DC electrical circuit, or both the hot and neutral sides of an AC circuit. This technique is used to prevent shock hazard in electric devices connected with unpolarised AC power plugs and sockets. Double switching is a crucial

safety engineering practice in railway signalling, wherein it is used to ensure that a single false feed of current to a relay is unlikely to cause a wrong side failure. It is an example of using redundancy to increase safety and reduce the likelihood of failure, analogous to double insulation. Double switching increases the cost and complexity of systems in which it is employed, for example by extra relay contacts and extra relays, so the technique is applied selectively where it can provide a cost-effective safety improvement.

## *Examples*

### Landslip and Washaway Detectors

A landslip or washaway detector is buried in the earth embankment, and opens a circuit should a landslide occur. It is not possible to guarantee that the wet earth of the embankment will not complete the circuit which is supposed to break. If the circuit is double cut with positive and negative wires, any wet conductive earth is likely to blow a fuse on the one hand, and short the detecting relay on the other hand, either of which is almost certain to apply the correct warning signal.

## *Accidents*

### Clapham

The Clapham Junction rail crash of 1988 was caused in part by the lack of double switching (known as "double cutting" in the British Railway industry). The signal relay in question was switched only on the hot side, while the return current came back on an unswitched wire. A loose wire bypassed the contacts by which the train detection relays switched the signal, allowing the signal to show green when in fact there was a stationary train ahead. 35 people were killed in the resultant collision.

### United Flight 811

A similar accident on the United Airlines Flight 811 was caused in part by a single-switched safety circuit for the baggage door mechanism. Failure of the wiring insulation in that circuit allowed the baggage door to be unlocked by a false feed, leading to a catastrophic de-pressurisation, and the deaths of nine passengers.

### Tri-Colour LED

Some tri-colour Light Emitting Diodes for railway use were wired with four wires, one for each of the three colours, and a common wire for the return. Due to water ingress and other problems, the lamp units were displaying false greens. The solution was to change to wiring with six wires with separate positive and negative wires to the LEDs of each colour.

### Faulty Attitude Indicator

Big airplanes have three independent attitude indicators, one for the pilot, one for the co-pilot, and a third one to resolve disputes between the first two. A Peruvian airplane apparently had a faulty wire in one of the indicators. The indicators for the pilot and co-pilot were switched to common mode, so they both displayed the same wrong attitude indications. In the dark, it was not possible to tell the true horizon in any way other than the attitude indicator, and the plane crashed into the sea.

### *Railway couplings*

Around 1994, new standards for the electrical couplings between carriages of United Kingdom passenger trains introduced the requirement for separate earth wires for critical functions such as brakes and doors. Common earths can cause interference between circuits that are otherwise independent, with unpredictable effects.

# Emergency power system


A backup generator for a large apartment building


A backup power fuel cell for telecom applications

**Emergency power systems** are a type of system, which may include lighting, generators, fuel cells and other apparatus, to provide backup power resources in a crisis or when regular systems fail. They find uses in a wide variety of settings from residential homes to hospitals, scientific laboratories, data centers, telecommunication equipment and modern naval ships. Emergency power systems can rely on generators, deep cycle batteries, flywheel energy storage or hydrogen fuel cells. Finally, some homebrew emergency power systems use regular lead-acid car batteries.

## History

Emergency power systems were used as early as World War II on naval ships. In combat, a ship may lose the function of its steam engines, which power the steam driven turbines for the generator. In such a case, one or more diesel engine(s) are used to drive back-up generators. Early transfer switches relied on manual operation; two switches would be placed horizontally, in line and the "on" position facing each other. a rod is placed in between. In order to operate the switch one source must be turned off, the rod moved to the other side and the other source turned on.

## Operation in buildings

Emergency power generator in a drinking water pumping station. Brons engine with Heemaf generator.


Another generator, powered by fossil fuels and used at a construction site

Mains power can be lost due to downed lines, malfunctions at a sub-station, inclement weather, planned blackouts or in extreme cases a grid-wide failure. In modern buildings, most emergency power systems have been and are still based on generators. Usually, these generators are diesel engine driven, although smaller buildings may use a gasoline engine driven generator and larger ones a gas turbine. However, lately, more use is being made of deep cycle batteries and other technologies such as flywheel energy storage or fuel cells. These latter systems do not produce polluting gases, thereby allowing the placement to be done within the building. Also, as a second advantage, they do not require a separate shed to be built for fuel storage.

With regular generators, an automatic transfer switch is used to connect emergency power. One side is connected to both the normal power feed and the emergency power feed; and the other side is connected to the load designated as emergency. If no electricity comes in on the normal side, the transfer switch uses a solenoid to throw a triple pole, single throw switch. This switches the feed from normal to emergency power. The loss of normal power also triggers a battery operated starter system to start the generator, similar to using a car battery to start an engine. Once the transfer switch is switched and the generator starts, the building's emergency power comes back on (after going off when normal power was lost.)

Unlike emergency lights, emergency lighting is not a type of light fixture; it is a pattern of the building's normal lights that provides a path of lights to allow for safe exit, or lights up service areas such as mechanical rooms and electric rooms. Exit signs, Fire alarm systems and the electric motor pumps for the fire sprinklers are almost always on emergency power. Other equipment on emergency power may include smoke isolation dampers, smoke evacuation fans, elevators, handicap doors and outlets in service areas. Hospitals use emergency power outlets to power life support systems and monitoring equipment. Some buildings may even use emergency power as part of normal operations, such as a theater using it to power show equipment because "the show must go on."

## *Operation in aviation*

Localizer, glideslope, and other instrument landing aids (such as microwave transmitters) are both high power consumers and mission-critical, and cannot be reliably operated from a battery supply, even for short periods. Hence, when absolute reliability is required (such as when Category 3 operations are in force at the airport) it is usual to run the system from a diesel generator with automatic switchover to the mains supply should the generator fail. This avoids any interruption to transmission while a generator is brought up to operating speed.

This is opposed to the typical view of emergency power systems, where the backup generators are seen as secondary to the mains electrical supply.

## *Electronic device protection*

Computers, communication networks and other modern electronic devices need not only power, but also a steady flow of it to continue to operate. If the source voltage drops significantly or drops out completely these devices will fail, even if it is for a fraction of a second. Because of this, even a generator back-up does not provide protection because of the start-up time involved.

To achieve this, extra equipment such as surge protectors, inverters, or a sometimes a complete uninterruptible power supply (UPS) is used. UPS systems can be local or building wide. A local UPS is a small box that fits under a desk or a telecom rack and powers a small number of devices. A building wide UPS can take on several different forms, depending on the application. It directly feeds a system of outlets designated as UPS feed and can power a large number of devices.

Since telephone exchanges use DC, the building's battery room is generally wired directly to the consuming equipment and floats continuously on the output of the rectifiers that normally supply DC rectified from utility power. When utility power fails, the battery carries the load without needing to switch. With this simple though somewhat expensive system, some exchanges have never lost power for a moment since the 1920s.

## *Structure and operation in utility stations*

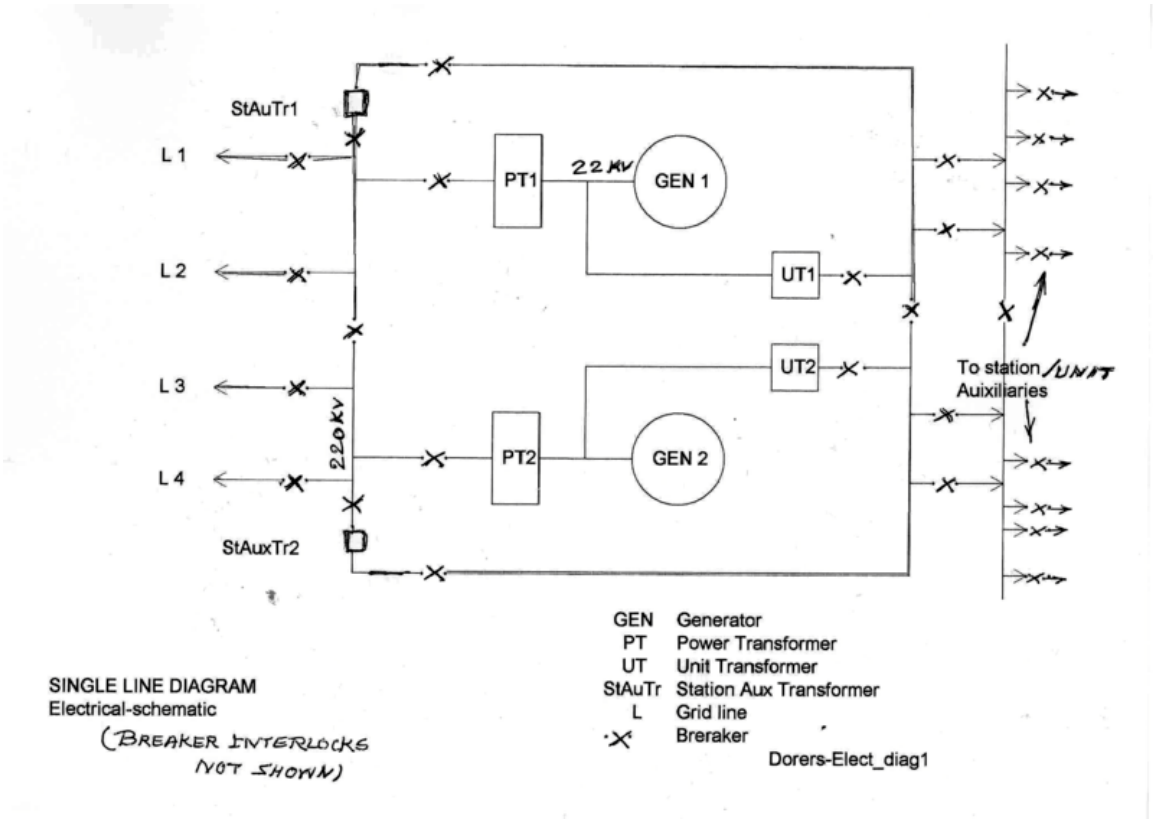


Diagram of a redundant power supply system.

In recent years, large units of a utility power station are usually designed on a unit system basis in which the required devices, including the boiler, the turbine generator unit, and its power (step up) and unit (auxiliary) transformer are solidly connected as one unit. A less common set-up consists of two units grouped together with one common station auxiliary. As each turbine generator unit has its own attached unit auxiliary transformer, it is connected to the circuit automatically. For starting the unit, the auxiliaries are supplied with power by another unit (auxiliary) transformer or station auxiliary transformer. The period of switching from the first unit transformer to the next unit is designed for automatic, instantaneous operation in times when the emergency power system needs to kick in. It is imperative that the power to unit auxiliaries not fail during a station shutdown (an occurrence known as black-out when all regular units temporarily fail). Instead, during shutdowns the grid is expected to remain operational. When problems occur, it is usually due to reverse power relays and frequency-operated relays on grid lines due to severe grid disturbances. Under these circumstances, the emergency station supply must kick in to avoid damage to any equipment and to prevent hazardous situations such as the release of hydrogen gas from generators to the local environment.

## *In nuclear power plants*

Emergency power systems, called there Emergency Diesel Generators (EDGs), are a required feature in nuclear power plants. They are typically installed in sets of three. The EDG installation is designed to the same safety-grade requirements as the other safety systems in the plant. The next (upcoming) generation of nuclear power plants includes some designs with multiple independent banks of EDGs (as in the ABWRs ).

## *Controlling the emergency power system*

For a 208 VAC emergency supply system, a central battery system with automatic controls, located in the power station building itself, is used to avoid long electric supply wires. This central battery system consists of lead-acid battery cell units to make up a 12 or 24 VDC system as well as stand-by cells, each with its own battery charging unit. Also needed are a voltage sensing unit capable of receiving 208 VAC and an automatic system that is able to signal to and activate the emergency supply circuit in case of failure of 208 VAC station supply.

## Chapter 7

# Error Detection and Correction

In information theory and coding theory with applications in computer science and telecommunication, **error detection and correction** or **error control** are techniques that enable reliable delivery of digital data over unreliable communication channels. Many communication channels are subject to channel noise, and thus errors may be introduced during transmission from the source to a receiver. Error detection techniques allow detecting such errors, while error correction enables reconstruction of the original data.

The general definitions of the terms are as follows:

- *Error detection* is the detection of errors caused by noise or other impairments during transmission from the transmitter to the receiver.
- *Error correction* is the detection of errors and reconstruction of the original, error-free data.

Error correction may generally be realized in two different ways:

- *Automatic repeat request (ARQ)* (sometimes also referred to as *backward error correction*): This is an error control technique whereby an error detection scheme is combined with requests for retransmission of erroneous data. Every block of data received is checked using the error detection code used, and if the check fails, retransmission of the data is requested – this may be done repeatedly, until the data can be verified.
- *Forward error correction (FEC)*: The sender encodes the data using an *error-correcting code (ECC)* prior to transmission. The additional information (redundancy) added by the code is used by the receiver to recover the original data. In general, the reconstructed data is what is deemed the "most likely" original data.

ARQ and FEC may be combined, such that minor errors are corrected without retransmission, and major errors are corrected via a request for retransmission: this is called *hybrid automatic repeat-request (HARQ)*.

## Introduction

The general idea for achieving error detection and correction is to add some redundancy (i.e., some extra data) to a message, which receivers can use to check consistency of the delivered message, and to recover data determined to be erroneous. Error-detection and correction schemes can be either systematic or non-systematic: In a systematic scheme, the transmitter sends the original data, and attaches a fixed number of *check bits* (or *parity data*), which are derived from the data bits by some deterministic algorithm. If only error detection is required, a receiver can simply apply the same algorithm to the received data bits and compare its output with the received check bits; if the values do not match, an error has occurred at some point during the transmission. In a system that uses a non-systematic code, the original message is transformed into an encoded message that has at least as many bits as the original message.

Good error control performance requires the scheme to be selected based on the characteristics of the communication channel. Common channel models include memory-less models where errors occur randomly and with a certain probability, and dynamic models where errors occur primarily in bursts. Consequently, error-detecting and correcting codes can be generally distinguished between *random-error-detecting/correcting* and *burst-error-detecting/correcting*. Some codes can also be suitable for a mixture of random errors and burst errors.

If the channel capacity cannot be determined, or is highly varying, an error-detection scheme may be combined with a system for retransmissions of erroneous data. This is known as automatic repeat request (ARQ), and is most notably used in the Internet. An alternate approach for error control is hybrid automatic repeat request (HARQ), which is a combination of ARQ and error-correction coding.

## Error detection schemes

Error detection is most commonly realized using a suitable hash function (or checksum algorithm). A hash function adds a fixed-length *tag* to a message, which enables receivers to verify the delivered message by recomputing the tag and comparing it with the one provided.

There exists a vast variety of different hash function designs. However, some are of particularly widespread use because of either their simplicity or their suitability for detecting certain kinds of errors (e.g., the cyclic redundancy check's performance in detecting burst errors).

Random-error-correcting codes based on minimum distance coding can provide a suitable alternative to hash functions when a strict guarantee on the minimum number of errors to be detected is desired. Repetition codes, described below, are special cases of error-correcting codes: although rather inefficient, they find applications for both error correction and detection due to their simplicity.

## Repetition codes

A **repetition code** is a coding scheme that repeats the bits across a channel to achieve error-free communication. Given a stream of data to be transmitted, the data is divided into blocks of bits. Each block is transmitted some predetermined number of times. For example, to send the bit pattern "1011", the four-bit block can be repeated three times, thus producing "1011 1011 1011". However, if this twelve-bit pattern was received as "1010 1011 1011" – where the first block is unlike the other two – it can be determined that an error has occurred.

Repetition codes are not very efficient, and can be susceptible to problems if the error occurs in exactly the same place for each group (e.g., "1010 1010 1010" in the previous example would be detected as correct). The advantage of repetition codes is that they are extremely simple, and are in fact used in some transmissions of numbers stations.

## Parity bits

A **parity bit** is a bit that is added to a group of source bits to ensure that the number of set bits (i.e., bits with value 1) in the outcome is even or odd. It is a very simple scheme that can be used to detect single or any other odd number (i.e., three, five, etc.) of errors in the output. An even number of flipped bits will make the parity bit appear correct even though the data is erroneous.

Extensions and variations on the parity bit mechanism are horizontal redundancy checks, vertical redundancy checks, and "double," "dual," or "diagonal" parity (used in RAID-DP).

## Checksums

A **checksum** of a message is a modular arithmetic sum of message code words of a fixed word length (e.g., byte values). The sum may be negated by means of a one's-complement prior to transmission to detect errors resulting in all-zero messages.

Checksum schemes include parity bits, check digits, and longitudinal redundancy checks. Some checksum schemes, such as the Luhn algorithm and the Verhoeff algorithm, are specifically designed to detect errors commonly introduced by humans in writing down or remembering identification numbers.

## Cyclic redundancy checks (CRCs)

A **cyclic redundancy check (CRC)** is a single-burst-error-detecting cyclic code and non-secure hash function designed to detect accidental changes to digital data in computer networks. It is characterized by specification of a so-called *generator polynomial*, which is used as the divisor in a polynomial long division over a finite field, taking the input data as the dividend, and where the remainder becomes the result.

Cyclic codes have favorable properties in that they are well suited for detecting burst errors. CRCs are particularly easy to implement in hardware, and are therefore commonly used in digital networks and storage devices such as hard disk drives.

Even parity is a special case of a cyclic redundancy check, where the single-bit CRC is generated by the divisor $x+1$.

### Cryptographic hash functions

A **cryptographic hash function** can provide strong assurances about data integrity, provided that changes of the data are only accidental (i.e., due to transmission errors). Any modification to the data will likely be detected through a mismatching hash value. Furthermore, given some hash value, it is infeasible to find some input data (other than the one given) that will yield the same hash value. Message authentication codes, also called *keyed* cryptographic hash functions, provide additional protection against intentional modification by an attacker.

### Error-correcting codes

Any error-correcting code can be used for error detection. A code with *minimum Hamming distance*, $d$, can detect up to $d-1$ errors in a code word. Using minimum-distance-based error-correcting codes for error detection can be suitable if a strict limit on the minimum number of errors to be detected is desired.

Codes with minimum Hamming distance $d=2$ are degenerate cases of error-correcting codes, and can be used to detect single errors. The parity bit is an example of a single-error-detecting code.

The Berger code is an early example of a unidirectional error(-correcting) code that can detect any number of errors on an asymmetric channel, provided that only transitions of cleared bits to set bits *or* set bits to cleared bits can occur.

## *Error correction*

### Automatic repeat request

Automatic Repeat reQuest (ARQ) is an error control method for data transmission that makes use of error-detection codes, acknowledgment and/or negative acknowledgment messages, and timeouts to achieve reliable data transmission. An *acknowledgment* is a message sent by the receiver to indicate that it has correctly received a data frame.

Usually, when the transmitter does not receive the acknowledgment before the timeout occurs (i.e., within a reasonable amount of time after sending the data frame), it retransmits the frame until it is either correctly received or the error persists beyond a predetermined number of retransmissions.

Three types of ARQ protocols are Stop-and-wait ARQ, Go-Back-N ARQ, and Selective Repeat ARQ.

ARQ is appropriate if the communication channel has varying or unknown capacity, such as is the case on the Internet. However, ARQ requires the availability of a back channel, results in possibly increased latency due to retransmissions, and requires the maintenance of buffers and timers for retransmissions, which in the case of network congestion can put a strain on the server and overall network capacity.

## Error-correcting code

An error-correcting code (ECC) or forward error correction (FEC) code is a system of adding redundant data, or *parity data*, to a message, such that it can be recovered by a receiver even when a number of errors (up to the capability of the code being used) were introduced, either during the process of transmission, or on storage. Since the receiver does not have to ask the sender for retransmission of the data, a back-channel is not required in forward error correction, and it is therefore suitable for simplex communication such as broadcasting. Error-correcting codes are frequently used in lower-layer communication, as well as for reliable storage in media such as CDs, DVDs, hard disks, and RAM.

Error-correcting codes are usually distinguished between convolutional codes and block codes:

- *Convolutional codes* are processed on a bit-by-bit basis. They are particularly suitable for implementation in hardware, and the Viterbi decoder allows optimal decoding.

- *Block codes* are processed on a block-by-block basis. Early examples of block codes are repetition codes, Hamming codes and multidimensional parity-check codes. They were followed by a number of efficient codes, Reed-Solomon codes being the most notable due to their current widespread use. Turbo codes and low-density parity-check codes (LDPC) are relatively new constructions that can provide almost optimal efficiency.

Shannon's theorem is an important theorem in forward error correction, and describes the maximum information rate at which reliable communication is possible over a channel that has a certain error probability or signal-to-noise ratio (SNR). This strict upper limit is expressed in terms of the channel capacity. More specifically, the theorem says that there exist codes such that with increasing encoding length the probability of error on a discrete memoryless channel can be made arbitrarily small, provided that the code rate is smaller than the channel capacity. The code rate is defined as the fraction $k/n$ of $k$ source symbols and $n$ encoded symbols.

The actual maximum code rate allowed depends on the error-correcting code used, and may be lower. This is because Shannon's proof was only of existential nature, and did not

show how to construct codes which are both optimal and have efficient encoding and decoding algorithms.

### Hybrid schemes

Hybrid ARQ is a combination of ARQ and forward error correction. There are two basic approaches:

- Messages are always transmitted with FEC parity data (and error-detection redundancy). A receiver decodes a message using the parity information, and requests retransmission using ARQ only if the parity data was not sufficient for successful decoding (identified through a failed integrity check).
- Messages are transmitted without parity data (only with error-detection information). If a receiver detects an error, it requests FEC information from the transmitter using ARQ, and uses it to reconstruct the original message.

The latter approach is particularly attractive on an erasure channel when using a rateless erasure code.

## *Applications*

Applications that require low latency (such as telephone conversations) cannot use Automatic Repeat reQuest (ARQ); they must use Forward Error Correction (FEC). By the time an ARQ system discovers an error and re-transmits it, the re-sent data will arrive too late to be any good.

Applications where the transmitter immediately forgets the information as soon as it is sent (such as most television cameras) cannot use ARQ; they must use FEC because when an error occurs, the original data is no longer available. (This is also why FEC is used in data storage systems such as RAID and distributed data store).

Applications that use ARQ must have a return channel. Applications that have no return channel cannot use ARQ.

Applications that require extremely low error rates (such as digital money transfers) must use ARQ.

### The Internet

In a typical TCP/IP stack, error control is performed at multiple levels:

- Each Ethernet frame carries a CRC-32 checksum. Frames received with incorrect checksums are discarded by the receiver hardware.
- The IPv4 header contains a checksum protecting the contents of the header. Packets with mismatching checksums are dropped within the network or at the receiver.

- The checksum was omitted from the IPv6 header in order to minimize processing costs in network routing and because current link layer technology is assumed to provide sufficient error detection.
- UDP has an optional checksum covering the payload and addressing information from the UDP and IP headers. Packets with incorrect checksums are discarded by the operating system network stack. The checksum is optional under IPv4, only, because the IP layer checksum may already provide the desired level of error protection.
- TCP provides a checksum for protecting the payload and addressing information from the TCP and IP headers. Packets with incorrect checksums are discarded within the network stack, and eventually get retransmitted using ARQ, either explicitly (such as through triple-ack) or implicitly due to a timeout.

## Deep-space telecommunications

Development of error-correction codes was tightly coupled with the history of deep-space missions due to the extreme dilution of signal power over interplanetary distances, and the limited power availability aboard space probes. Whereas early missions sent their data uncoded, starting from 1968 digital error correction was implemented in the form of (sub-optimally decoded) convolutional codes and Reed-Muller codes. The Reed-Muller code was well suited to the noise the spacecraft was subject to (approximately matching a bell curve), and was implemented at the Mariner spacecraft for missions between 1969 and 1977.

The Voyager 1 and Voyager 2 missions, which started in 1977, were designed to deliver color imaging amongst scientific information of Jupiter and Saturn. This resulted in increased coding requirements, and thus the spacecraft were supported by (optimally Viterbi-decoded) convolutional codes that could be concatenated with an outer Golay (24,12,8) code. The Voyager 2 probe additionally supported an implementation of a Reed-Solomon code: the concatenated Reed-Solomon-Viterbi (RSV) code allowed for very powerful error correction, and enabled the spacecraft's extended journey to Uranus and Neptune.

The CCSDS currently recommends usage of error correction codes with performance similar to the Voyager 2 RSV code as a minimum. Concatenated codes are increasingly falling out of favor with space missions, and are replaced by more powerful codes such as Turbo codes or LDPC codes.

The different kinds of deep space and orbital missions that are conducted suggest that trying to find a "one size fits all" error correction system will be an ongoing problem for some time to come. For missions close to earth the nature of the channel noise is different from that of a spacecraft on an interplanetary mission experiences. Additionally, as a spacecraft increases its distance from earth, the problem of correcting for noise gets larger.

## Satellite broadcasting (DVB)

The demand for satellite transponder bandwidth continues to grow, fueled by the desire to deliver television (including new channels and High Definition TV) and IP data. Transponder availability and bandwidth constraints have limited this growth, because transponder capacity is determined by the selected modulation scheme and Forward error correction (FEC) rate.

Overview

- QPSK coupled with traditional Reed Solomon and Viterbi codes have been used for nearly 20 years for the delivery of digital satellite TV.
- Higher order modulation schemes such as 8PSK, 16QAM and 32QAM have enabled the satellite industry to increase transponder efficiency by several orders of magnitude.
- This increase in the information rate in a transponder comes at the expense of an increase in the carrier power to meet the threshold requirement for existing antennas.
- Tests conducted using the latest chipsets demonstrate that the performance achieved by using Turbo Codes may be even lower than the 0.8 dB figure assumed in early designs.

## Data storage

Error detection and correction codes are often used to improve the reliability of data storage media.

A "parity track" was present on the first magnetic tape data storage in 1951. The "Optimal Rectangular Code" used in group code recording tapes not only detects but also corrects single-bit errors.

Some file formats, particularly archive formats, include a checksum (most often CRC32) to detect corruption and truncation and can employ redundancy and/or parity files to recover portions of corrupted data.

Reed Solomon codes are used in compact discs to correct errors caused by scratches.

Modern hard drives use CRC codes to detect and Reed-Solomon codes to correct minor errors in sector reads, and to recover data from sectors that have "gone bad" and store that data in the spare sectors.

RAID systems use a variety of error correction techniques, to correct errors when a hard drive completely fails.

## Error-correcting memory

DRAM memory may provide increased protection against soft errors by relying on error correcting codes. Such error-correcting memory, known as *ECC* or *EDAC-protected* memory, is particularly desirable for high fault-tolerant applications, such as servers, as well as deep-space applications due to increased radiation.

Error-correcting memory controllers traditionally use Hamming codes, although some use triple modular redundancy.

Interleaving allows distributing the effect of a single cosmic ray potentially upsetting multiple physically neighboring bits across multiple words by associating neighboring bits to different words. As long as a single event upset (SEU) does not exceed the error threshold (e.g., a single error) in any particular word between accesses, it can be corrected (e.g., by a single-bit error correcting code), and the illusion of an error-free memory system may be maintained.

# Chapter 8

# Fail-Safe

A **fail-safe** or **fail-secure** device is one that, in the event of failure, responds in a way that will cause no harm, or at least a minimum of harm, to other devices or danger to personnel.

Fail-safe components should not be confused with fail-secure components. A fail-secure component will allow, but does not cause, a system failure. For example, a fail-secure lock will remain locked during a failure and cannot be unlocked, even with the correct key. In contrast, a fail-safe component does not allow a system failure. For example, a lock that unlocks at the wrong time has failed, but is considered fail-safe because it does not open or attract undue attention to the door's unlocked state.

Significantly, despite popular belief to the contrary, "fail-safe" does not mean that failure is improbable, but rather that a system's design mitigates any unsafe consequences of failure.

## *Examples*

### Mechanical or physical

An aircraft lights its afterburners to maintain full power following an arrested landing aboard an aircraft carrier.

- Aircraft landing on an aircraft carrier increases the throttle to full power at touchdown. If the arresting wires fail to capture the plane, it is able to take off again.
- Coiling/rolling fire doors that are activated by building alarm systems or local smoke detectors must close automatically when signaled regardless of power. In case of power outage the coiling fire door does not need to close, but must be capable of automatic closing when given a signal from the building alarm systems or smoke detectors. A temperature sensitive fusible link may be employed to hold the fire doors open against gravity or a closing spring. In case of fire, the link melts and releases the doors, and they close.
- Luggage carts in airports in which the hand-brake must be held down at all times. If it is released, the cart will stop.
- Lawnmowers and snow blowers have a hand-closed lever that must be held down at all times. If it is released, it stops the blade's or rotor's rotation.
- Air brakes on railway trains and air brakes on trucks. The brakes are held in the 'off' position by air pressure created in the brake system. Should a brake line split, or a carriage become de-coupled, the air pressure will be lost and the brakes

applied. It is impossible to drive a train or truck with a serious leak in the air brake system.
- Motorized gates – In case of power outage the gate can be pushed open by hand with no crank or key required. However, as this would allow virtually anyone to go through the gate, a *fail-secure* design is used: In a power outage, the gate can only be opened by a hand crank that is usually kept in a safe area.
- During early Apollo program missions to the Moon, the spacecraft was put on a free return trajectory – if the engines failed at lunar orbit insertion, the craft would safely coast back to Earth.
- Elevator cabins that begin to accelerate too quickly as a result of the cables failing have a safety mechanism which uses contact with the guide rail to decelerate the car.
- Various devices that operate with fluids use fuses or valves as a fail-safe mechanism.
- A railway semaphore signal of the upper quadrant type is designed so that should the signal arm be weighed down by snow or the cable controlling the signal break the arm, returns to the 'danger' position, preventing any trains passing the inoperative signal.
- Diving watches – On diving watches the bezel is "unidirectional", i.e., it contains a ratchet so it can only be turned anti-clockwise to increase the apparent elapsed time. If the bezel could be turned the other way this could suggest to a diver that the elapsed time was shorter than the truth, thus giving a falsely low elapsed time reading and therefore an assumed falsely low air consumption reading and falsely high remaining air reading, all of which could be highly dangerous. In this fashion, the one-way bezel is designed to be "fail-safe", that is, if it 'fails', i.e., by being inadvertently rotated during the dive, it will only rotate so as to give a false reading of increased time below and thus less assumed tank air remaining rather than the opposite; it 'fails' in a 'safe' way.

## Electrical or electronic

- Many devices are protected from short circuit with fuses. The destruction of the fuse will prevent destruction of the device.
- Avionics using redundant systems to perform the same computation with voting logic to determine the "safe" result.
- Traffic light controllers use a *Conflict Monitor Unit* to detect faults or conflicting signals and switch an intersection to all flashing red, rather than displaying potentially dangerous conflicting signals, e.g. showing green in all directions.
- The automatic protection of programs and/or processing systems when a hardware or software failure is detected in a computer system. A classic example is a watchdog timer.
- A control operation or function that prevents improper system functioning or catastrophic degradation in the event of circuit malfunction or operator error; for example, the **failsafe** track circuit used to control railway block signals.

- The iron pellet ballast on the Bathyscaphe is dropped to allow the submarine to ascend. The ballast is held in place by electromagnets. If electrical power fails the ballast is released, and the submarine then ascends to safety.
- Inside a modern CPU are features to prevent damage through overheating. In the event of cooling failure, the CPU will throttle then shut down beyond a critical temperature threshold to avoid damage.
- In industrial automation, alarm signals are usually "normally closed" (or active at 0). This insures that in case of a wire break the alarm will be triggered. If the signal were normally open, no wire failure would be detected.

### Procedural

As well as physical devices and systems fail-safe procedures can be created so that if a procedure is not carried out or carried out incorrectly no dangerous action results. For example:

- In railway signalling signals which are not in active use for a train are required to be kept in the 'danger' position. The default position of every signal is therefore 'danger,' and therefore a positive action—setting signals to 'clear'—is required before a train may pass. This practice also ensures that, in case of a fault in the signalling system, an incapacitated signalman, or the unexpected entry of a train, that a train will never be shown an erroneous 'clear' signal.
- Train drivers are instructed that a railway signal showing a confusing, contradictory or unfamiliar aspect (for example a colour light signal that has suffered an electrical failure and is showing no light at all) must be treated as showing 'danger'. In this way, the driver contributes to the fail-safety of the system.

## *Other terminology*

Fail-safe (foolproof) devices are also known as *poka-yoke* devices. *Poka-yoke*, a Japanese term, was coined by Shigeo Shingo, a quality guru.

# Chapter 9

# Fly-by-Wire

**Fly-by-wire**


Green colored flight control wiring of a test aircraft

A **fly-by-wire** (FBW) system replaces manual flight control of an aircraft with an electronic interface. The movements of flight controls are converted to electronic signals transmitted by wires (hence the fly-by-wire term), and flight control computers determine how to move the actuators at each control surface to provide the ordered response. The fly-by-wire system also allows automatic signals sent by the aircraft's computers to perform functions without the pilot's input, as in systems that automatically help stabilize the aircraft.

## *Development*

Mechanical and hydro-mechanical flight control systems are relatively heavy and require careful routing of flight control cables through the aircraft by systems of pulleys, cranks, tension cables and hydraulic pipes. Both systems often require redundant backup to deal

with failures, which again increases weight. Furthermore, both have limited ability to compensate for changing aerodynamic conditions. Dangerous characteristics such as stalling, spinning and pilot-induced oscillation (PIO), which depend mainly on the stability and structure of the aircraft concerned rather than the control system itself, can still occur with these systems.

The term "fly-by-wire" implies a purely electrically-signaled control system. However, it is used in the general sense of computer-configured controls, where a computer system is interposed between the operator and the final control actuators or surfaces. This modifies the manual inputs of the pilot in accordance with control parameters.

Side-sticks, center sticks, or conventional flight control yokes can be used to fly FBW aircraft. While the side-stick offers the advantages of being lighter, mechanically simpler, and unobtrusive, The Boeing Company's aerospace engineers decided that the lack of visual feedback (none given by side-sticks) is a significant problem, and so they designed conventional control yokes in the Boeing 777 and the brand-new Boeing 787, which is undergoing flight tests as of June 2010. This same approach has been used for the Embraer 170/190 jets. Most Airbus airliners are operated with side-sticks.

## Basic operation

## Command

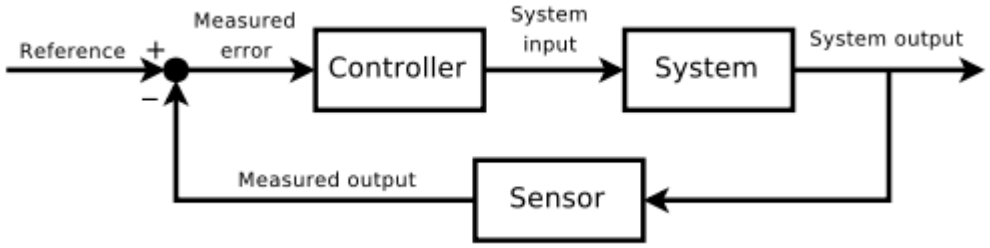


Simple feedback loop

Fly-by wire systems are by their nature quite complex however their operation can be explained in relatively simple terms. When a pilot moves the control column (or sidestick) a signal is sent to a computer, this is analogous to moving a game controller, the signal is sent through multiple wires (channels) to ensure that the signal reaches the computer. When there are three channels being used this is known as 'Triplex'. The computer receives the signals, performs a calculation (adds the signal voltages and divides by the number of signals received to find the mean average voltage) and adds another channel. These four 'Quadruplex' signals are then sent to the control surface actuator and the surface begins to move. Potentiometers in the actuator send a signal back to the computer (usually a negative voltage) reporting the position of the actuator. When the actuator reaches the desired position the two signals (incoming and outgoing) cancel each other out and the actuator stops moving (completing a feedback loop).

## Automatic Stability Systems

Fly-by-wire control systems allow aircraft computers to perform tasks without pilot input. Automatic stability systems operate in this way. Gyroscopes fitted with sensors are mounted in an aircraft to sense movement changes in the pitch, roll and yaw axes. Any movement (from straight and level flight for example) results in signals to the computer, which automatically moves control actuators to stabilize the aircraft.

## Safety and redundancy

Aircraft systems may be quadruplexed (four independent channels) to prevent loss of signals in the case of failure of one or even two channels. High performance aircraft that have FBW controls (also called CCVs or Control-Configured Vehicles) may be deliberately designed to have low or even negative aerodynamic stability in some flight regimes, the rapid-reacting CCV controls compensating for the lack of natural stability.

Pre-flight safety checks of a fly-by-wire system are often performed using Built-In Test Equipment (BITE). On programming the system, either by the pilot or groundcrew, a number of control movement steps are automatically performed. Any failure will be indicated to the crews.

Some aircraft, the Panavia Tornado for example, retain a very basic hydro-mechanical backup system for limited flight control capability on losing electrical power, in the case of the Tornado this allows rudimentary control of the stabilators only for pitch and roll axis movements.

## Weight Saving

A FBW aircraft can be lighter than a similar design with conventional controls. Partly due to the lower overall weight of the system components; and partly because the natural aerodynamic stability of the aircraft can be relaxed, slightly for a transport aircraft and more for a maneuverable fighter, which means that the stability surfaces that are part of the aircraft structure can therefore be made smaller. These include the vertical and horizontal stabilizers (fin and tailplane) that are (normally) at the rear of the fuselage. If these structures can be reduced in size, airframe weight is reduced. The advantages of FBW controls were first exploited by the military and then in the commercial airline market. The Airbus series of airliners used full-authority FBW controls beginning with their A320 series. Boeing followed with their 777 and later designs.

Electronic fly-by-wire systems can respond flexibly to changing aerodynamic conditions, by tailoring flight control surface movements so that aircraft response to control inputs is appropriate to flight conditions. Electronic systems require less maintenance, whereas mechanical and hydraulic systems require lubrication, tension adjustments, leak checks, fluid changes, etc. Furthermore, putting circuitry between pilot and aircraft can enhance safety; for example the control system can try to prevent a stall, or it can stop the pilot from over stressing the airframe.

The main concern with fly-by-wire systems is reliability. While traditional mechanical or hydraulic control systems usually fail gradually, the loss of all flight control computers could immediately render the aircraft uncontrollable. For this reason, most fly-by-wire systems incorporate either redundant computers (triplex, quadruplex etc.), some kind of mechanical or hydraulic backup or a combination of both. A "mixed" control system such as the latter is not desirable and modern FBW aircraft normally avoid it by having more independent FBW channels, thereby reducing the possibility of overall failure to minuscule levels that are acceptable to the independent regulatory and safety authority responsible for aircraft design, testing and certification before operational service.

## History


F-8C Crusader digital fly-by-wire testbed

Electronic signalling of the control surfaces was tested in the 1950s. This replaced long runs of mechanical and hydraulic connections with electrical ones.

The first non-experimental aircraft that was designed and flown (in 1958) with a fly-by-wire flight control system was the Avro Canada CF-105 Arrow. a feat not repeated with a production aircraft until Concorde in 1969. This system also included solid-state components and system redundancy, was designed to be integrated with a computerised navigation and automatic search and track radar, was flyable from ground control with data uplink and downlink, and provided artificial feel (feedback) to the pilot.

The first digital fly-by-wire aircraft to take to the air (in 1972) was an F-8 Crusader, which had been modified electronically by the National Aeronautics and Space Administration of the United States as a test aircraft, a feat mirrored in the USSR by the Sukhoi T-4. At about the same time in the United Kingdom a trainer variant of the British

Hawker Hunter fighter was modified at the British Royal Aircraft Establishment with fly-by-wire flight controls for the right-seat pilot. This was test-flown, with the left-seat pilot having conventional flight controls for safety reasons, and with the capability for him to override and turn off the fly-by-wire system.

## *Analog systems*

All "fly-by-wire" flight control systems eliminate the complexity, the fragility, and the weight of the mechanical circuit of the hydromechanical or electromechanical flight control systems. Fly-by-wire replace those with electronic circuits. The control mechanisms in the cockpit now operate signal transducers, which in turn generate the appropriate electronic commands. These are next processed by an electronic controller, either an analog one, or more modernly, a digital one. Aircraft and spacecraft autopilots are now part of the electronic controller.

The hydraulic circuits are similar except that mechanical servo valves are replaced with electrically-controlled servo valves, operated by the electronic controller. This is the simplest and earliest configuration of an analog fly-by-wire flight control system. In this configuration, the flight control systems must simulate "feel". The electronic controller controls electrical feel devices that provide the appropriate "feel" forces on the manual controls. This was used in Concorde, the first production fly-by-wire airliner.

In more sophisticated versions, analog computers replaced the electronic controller. The canceled 1950s Canadian supersonic intercepter, the Avro Canada CF-105 Arrow, employed this type of system. Analog computers also allowed some customization of flight control characteristics, including relaxed stability. This was exploited by the early versions of F-16, giving it impressive maneuverability.

## *Digital systems*


The Airbus A320, first airliner with digital fly-by-wire controls

A digital fly-by-wire flight control system is similar to its analog counterpart. However, the signal processing is done by digital computers and the pilot literally can "fly-via-computer". This also increases the flexibility of the flight control system, since the digital computers can receive input from any aircraft sensor (such as the altimeters and the pitot tubes. This also increases the electronic stability, because the system is less dependent on the values of critical electrical components in an analog controller.

The computers sense position and force inputs from pilot controls and aircraft sensors. They solve differential equations to determine the appropriate command signals that move the flight controls to execute the intentions of the pilot.

The programming of the digital computers enable flight envelope protection. In this aircraft designers precisely tailor an aircraft's handling characteristics, to stay within the overall limits of what is possible given the aerodynamics and structure of the aircraft. For example, the computer in flight envelope protection mode can try to prevent the aircraft from being handled dangerously by preventing pilots from exceeding preset limits on the aircraft's flight-control envelope, such as those that prevent stalls and spins, and which limit airspeeds and g forces on the airplane. Software can also be included that stabilize the flight-control inputs to avoid pilot-induced oscillations.

Since the flight-control computers continuously "fly" the aircraft, pilot's workloads can be reduced. Also, in military and naval applications, it is now possible to fly military aircraft that have relaxed stability. The primary benefit for such aircraft is more maneuverability during combat and training flights, and the so-called "carefree handling" because stalling, spinning. and other undesirable performances are prevented automatically by the computers.

Digital flight control systems enable inherently unstable combat aircraft, such as the F-117 Nighthawk and the B-2 Spirit flying wing to fly in usable and safe manners.

## Applications

A Dassault Falcon 7X, the first business jet with digital fly-by-wire controls

- The Space Shuttle Orbiter has an all-digital fly-by-wire control system. This system was first exercised (as the only flight control system) during the glider unpowered-flight "Approach and Landing Tests" that began on the Space Shuttle *Enterprise* during 1977.
- During 1984, the Airbus Industries Airbus A320 became the first airliner to fly with an all-digital fly-by-wire control system.
- During 2005, the Dassault Falcon 7X became the first business jet with fly-by-wire controls.

## Legislation

The Federal Aviation Administration (FAA) of the United States has adopted the RTCA/DO-178B, titled "Software Considerations in Airborne Systems and Equipment Certification", as the certification standard for aviation software. Any safety-critical component in a digital fly-by-wire system including applications of the laws of aeronautics and computer operating systems will need to be certified to DO-178B Level A, which is applicable for preventing potential catastrophic failures.

Nevertheless, the top concern for computerized, digital, fly-by-wire systems is reliability, even more so than for analog electronic control systems. This is because the digital computers that are running software are often the only control path between the pilot and

aircraft's flight control surfaces. If the computer software crashes for any reason, the pilot may be unable to control an aircraft. Hence virtually all fly-by-wire flight control systems are either triply or quadruply redundant in their computers and electronics. These have three or four flight-control computers operating in parallel, and three or four separate data buses connecting them with each control surface.

## Redundancy

If one of the flight-control computers crashes, or is damaged in combat, or suffers from "insanity" caused by electromagnetic pulses, the others overrule the faulty one (or even two of them), they continue flying the aircraft safely, and they can either turn off or re-boot the faulty computers. Any flight-control computer whose results disagree with the others is ruled to be faulty, and it is either ignored or re-booted. (In other words, it is voted-out of control by the others.)

In addition, most of the early digital fly-by-wire aircraft also had an analog electrical, a mechanical, or a hydraulic back-up flight control system. The Space Shuttle has, in addition to its redundant set of four digital computers running its primary flight-control software, a fifth back-up computer running a separately developed, reduced-function, software flight-control system - one that can be commanded to take over in the event that a fault ever affects all of the computers in the other four. This back-up system serves to reduce the risk of total flight-control-system failure ever happening because of a general-purpose flight software fault has escaped notice in the other four computers.

For airliners, flight-control redundancy improves their safety, but fly-by-wire control systems also improve economy in flight because they are lighter, and they eliminate the need for many mechanical, and heavy, flight-control mechanisms. Furthermore, most modern airliners have computerized systems that control their jet engine throttles, air inlets, fuel storage and distribution system, in such a way to minimize their consumption of jet fuel. Thus, digital control systems do their best to reduce the cost of flights.

## Airbus/Boeing

Airbus and Boeing commercial airplanes differ in their approaches in using fly-by-wire systems. In Airbus airliners, the flight-envelope control system always retains ultimate flight control, and it will not permit the pilots to fly outside these performance limits. However, in the event of multiple failures of redundant computers, the A320 does have mechanical back-up system for its pitch trim and its rudder. The A340-600 has a purely electrical (not electronic) back-up rudder control system, and beginning with the new A380 airliner, all flight-control systems have back-up systems that are purely electrical through the use of a so-called "three-axis Backup Control Module" (BCM)

With the Boeing 777 model airliners, the two pilots can completely override the computerized flight-control system to permit the aircraft to be flown beyond its usual flight-control envelope during emergencies. Airbus's strategy, which began with the Airbus A320, has been continued on subsequent Airbus airliners.

## *Engine digital control*

The advent of FADEC (Full Authority Digital Engine Control) engines permits operation of the flight control systems and autothrottles for the engines to be fully integrated. On modern military aircraft other systems such as autostabilization, navigation, radar and weapons system are all integrated with the flight control systems. FADEC allows maximum performance to be extracted from the aircraft without fear of engine misoperation, aircraft damage or high pilot workloads. In the civil field, the integration increases flight safety and economy. The Airbus A320 and its fly-by-wire brethren are protected from dangerous situations such as low-speed stall or overstressing by flight envelope protection. As a result, in such conditions, the flight control systems commands the engines to increase thrust without pilot intervention. In economy cruise modes, the flight control systems adjust the throttles and fuel tank selections more precisely than all but the most skillful pilots. FADEC reduces rudder drag needed to compensate for sideways flight from unbalanced engine thrust. On the A330/A340 family, fuel is transferred between the main (wing and center fuselage) tanks and a fuel tank in the horizontal stabilizer, to optimize the aircraft's center of gravity during cruise flight. The fuel management controls keep the aircraft's center of gravity accurately trimmed with fuel weight, rather than drag-inducing aerodynamic trims in the elevators.

## *Further developments*

### Fly-by-optics

Fly-by-optics is sometimes used instead of fly-by-wire because it can transfer data at higher speeds, and it is immune to electromagnetic interference. In most cases, the cables are just changed from electrical to optical fiber cables. Sometimes it is referred to as "fly-by-light" due to its use of fiber optics. The data generated by the software and interpreted by the controller remain the same.

### Power-by-wire

Having eliminated the mechanical transmission circuits in fly-by-wire flight control systems, the next step is to eliminate the bulky and heavy hydraulic circuits. The hydraulic circuit is replaced by an electrical power circuit. The power circuits power electrical or self-contained electrohydraulic actuators that are controlled by the digital flight control computers. All benefits of digital fly-by-wire are retained.

The biggest benefits are weight savings, the possibility of redundant power circuits and tighter integration between the aircraft flight control systems and its avionics systems. The absence of hydraulics greatly reduces maintenance costs. This system is used in the Lockheed Martin F-35 Lightning II and in Airbus A380 backup flight controls. The Boeing 787 will also incorporate some electrically operated flight controls (spoilers and horizontal stabilizer), which will remain operational with either a total hydraulics failure and/or flight control computer failure.

## Fly-by-wireless

Wiring adds a considerable amount of weight to an aircraft; therefore, researchers are exploring implementing fly-by-wireless solutions. Fly-by-wireless systems are very similar to fly-by-wire systems, however, instead of using a wired protocol for the physical layer a wireless protocol is employed.

In addition to reducing weight, implementing a wireless solution has the potential to reduce costs throughout an aircraft's life cycle. For example, many key failure points associated with wire and connectors will be eliminated thus hours spent troubleshooting wires and connectors will be reduced. Furthermore, engineering costs could potentially decrease because less time would be spent on designing wiring installations, late changes in an aircraft's design would be easier to manage, etc.

## Intelligent Flight Control System

A newer flight control system, called Intelligent Flight Control System (IFCS), is an extension of modern digital fly-by-wire flight control systems. The aim is to intelligently compensate for aircraft damage and failure during flight, such as automatically using engine thrust and other avionics to compensate for severe failures such as loss of hydraulics, loss of rudder, loss of ailerons, loss of an engine, etc. Several demonstrations were made on a flight simulator where a Cessna-trained small-aircraft pilot successfully landed a heavily-damaged full-size concept jet, without prior experience with large-body jet aircraft. This development is being spearheaded by NASA Dryden Flight Research Center. It is reported that enhancements are mostly software upgrades to existing fully computerized digital fly-by-wire flight control systems.

# Chapter 10

# Hot Spare and Hot Swapping

## Hot spare

A **hot spare** or **hot standby** is used as a failover mechanism to provide reliability in system configurations. The hot spare is active and connected as part of a working system. When a key component fails, the hot spare is switched into operation. More generally, a hot standby can be used to refer to any device or system that is held in readiness to overcome an otherwise significant start-up delay.

### *Examples*

Examples of hot spares are components such as A/V switches, computers, network printers, and hard disks. The equipment is powered on, or considered "hot," but not actively functioning in (i.e. used by) the system.

Electrical generators may be held on hot standby, or a steam train may be held at the shed fired up (literally hot) ready to replace a possible failure of an engine in service.

### *Explanation*

In designing a reliable system, it is recognized that there will be failures. At the extreme, a complete system can be duplicated and kept up to date—so in the event of the primary system failing, the secondary system can be switched in with little or no interruption. More often, a hot spare is a single vital component without which the entire system would fail. The spare component is integrated into the system in such a way that in the event of a problem, the system can be altered to use the spare component. This may be done automatically or manually, but in either case it is normal to have some means of error detection. A hot spare does not necessarily give 100% availability or protect against temporary loss of the system during the switching process; it is designed to significantly reduce the time that the system is unavailable.

Hot standby may have a slightly different connotation of being active but not productive to hot spare, that is it is a state rather than object. For example, in a national power grid, the supply of power needs to be balanced to demand over a short term. It can take many hours to bring a coal-fired power station up to productive temperatures. To allow for load balancing, generator turbines may be kept running with the generators switched off so as peaks of demand occur, the generators can rapidly be switched on to balance the load. Being in the state of being ready to run is known as hot standby. Though it is not a modern phenomenon, steam train operators might hold a spare steam engine at a terminus fired up, as starting an engine cold would take a significant amount of time.

The spare may be similar component or system, or it may be a system of reduced performance, designed to cope for the duration of the time to repair and recover the original component. In high availability systems, it is common to design so that not only is there a spare that can quickly be switched in, but also that the failed component can be repaired or replaced without stopping the system - this is known as hot swapping. It may be considered that the probability of a second failure is low, and therefore the system is designed simply to allow operation to continue until a suitable maintenance period. The appropriate solution is normally determined by balancing the costs of implementing the availability against the likelihood of a problem and the severity of that problem.

## *Computer usage*

A **hot spare disk** is a disk or group of disks used to automatically or manually, depending upon the hot spare policy, replace a failing or failed disk in a RAID configuration. The hot spare disk reduces the mean time to recovery (MTTR) for the RAID redundancy group, thus reducing the probability of a second disk failure and the resultant data loss that would occur in any singly redundant RAID (e.g., RAID-1, RAID-5, RAID-10). Typically, a hot spare is available to replace a number of different disks and systems employing a hot spare normally require a redundant group to allow time for the data to be generated onto the spare disk. During this time the system is exposed to data loss due to a subsequent failure, and therefore the automatic switching to a spare disk reduces the time of exposure to that risk compared to manual discovery and implementation.

The concept of hot spares is not limited to hardware, but also software systems can be held in a state of readiness, for example a database server may have a software copy on hot standby, possibly even on the same machine to cope with the various factors that make a database unreliable, such as the impact of disc failure, poorly written queries or database software errors.

## *Hot standby operation in railway signalling*

At least two units of the same type will be powered up, receiving the same set of inputs, performing identical computations and producing identical outputs in a nearly-synchronous manner. The outputs are typically physical outputs (individual ON/OFF type digital signals, or analog signals), or serial data messages wrapped in suitable protocols

depending upon the nature of their intended use. Outputs from only one unit (designated as the master or on-line unit, via application logic) are used to control external devices (such as switches, signals, on-board propulsion/braking control devices, etc.) or simply to provide displays. The other unit is a hot-standby or a hot spare unit, ready to take over if the master unit fails. When the master unit fails, an automatic failover to the hot spare occurs within a very short time and the outputs from the hot spare, now the master unit, are delivered to the controlled devices and displays. The controlled devices and displays may experience a short blip or disturbance during the failover time. However, they can be designed to tolerate/ignore the disturbances so that the overall system operation is not affected.

# Hot swapping

**Hot swapping** and **hot plugging** are terms used to describe the functions of replacing computer system components without shutting down the system. More specifically, hot swapping describes replacing components without significant interruption to the system, while hot plugging describes the addition of components that would expand the system without significant interruption to the operation of the system. Once the appropriate software is installed on the computer, a user can plug and unplug the component without rebooting. A well-known example of this functionality is the Universal Serial Bus (USB) that allows users to add or remove peripheral components such as a mouse, keyboard, or printer.

## *Reasons for hot-swapping*

Hot swapping is used whenever it is desirable to change the configuration or repair a working system without interrupting its operation. It may simply be for convenience, to avoid the delay and nuisance of shutting down and then restarting complex equipment, or because it is essential that the equipment be permanently available.

Hot swapping may be used to add or remove peripherals or components, to allow a device to synchronize data with a computer, and to replace faulty modules without interrupting equipment operation.

Equipment may be designed with redundancy so that in the event of the failure of a component, other parts of the system carry out its functions while the faulty component is removed and a replacement connected. For example, computer RAID disk arrays allow a faulty disk to be hot-swapped for a new one; the new one is configured to become part of the array automatically or by user command. A machine may have dual power supplies, each adequate to power the machine; a faulty one may be hot-swapped.

Hot swapping can also be used as a means for circumventing security measures preventing the execution of unsigned code, a common example of this is hot swapping the Microsoft Xbox Console hard drive.

## *System considerations*

Machines that support hot swapping need to be able to modify their operation for the changed configuration, either automatically on detecting the change, or by user intervention. All electrical and mechanical connections associated with hot-swapping must be designed so that neither the equipment nor the user can be harmed while hot-swapping. Other components in the system must be designed so that the removal of a hot-swappable component does not interrupt operation.

Some implementations require a component shutdown procedure prior to removal. This simplifies the design, but such devices are not robust in the case of component failure. If a component is removed while it is being used, the operations to that device fail and the user is responsible for retrying if necessary, although this is not usually considered to be a problem.

More complex implementations may recommend but do not require that the component be shut down, with sufficient redundancy in the system to allow operation to continue if a component is removed without being shut down. In these systems hot swap is normally used for regular maintenance to the computer, or to replace a broken component.

There are two slightly differing meanings of the term *hot swapping*. It may refer only to the ability to add or remove hardware without powering down the system, while the system software may have to be notified by the user of the event in order to cope with it. Examples include RS-232 and lower-end SCSI devices. This is sometimes called cold plugging. However, if the system can detect and respond to addition or removal of hardware, it is referred to as *true hot plugging*. Examples include USB, FireWire and higher-end SCSI devices.

## *Connectors*

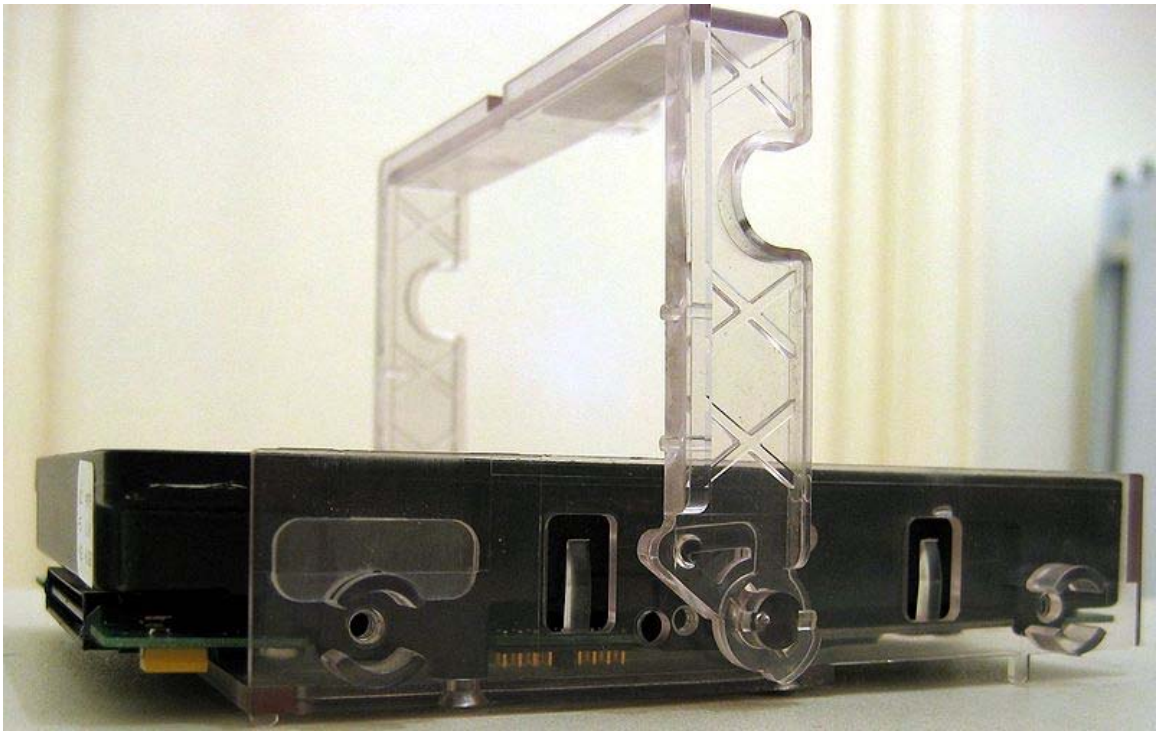

Sun SPARCstation hot swappable Single Connector Attachment (SCA) drive cradle

Most modern hot-swap methods use a specialized connector with staggered pins, so that certain pins are certain to be connected before others. At one time staggered pins were thought to be an expensive solution, but many contemporary connector families now come with staggered pins as standard; for example, they are used on all modern serial SCSI disk-drives. Specialized hot-plug power connector pins are now commercially available with repeatable DC current interruption ratings of up to 16 A. Printed circuit boards are made with staggered edge-fingers for direct hot-plugging into a backplane connector.

Most staggered-pin designs have ground pins longer than the others, ensuring that no sensitive circuitry is connected before there is a reliable system ground. The other pins may all be the same length; in some cases three pin lengths are used.

Although the speed of plugging cannot be controlled precisely, practical considerations will provide limits that can be used to determine worst-case conditions. For a typical staggered pin design where the length difference is 0.5 mm, the elapsed time between long and short pin contact is between 25 ms and 250 ms. It is quite practical to design hot-swap circuits that can operate over that dynamic range. Pins of the same nominal length do not necessarily make contact at exactly the same time due to mechanical tolerances, and angling of the connector when inserted.

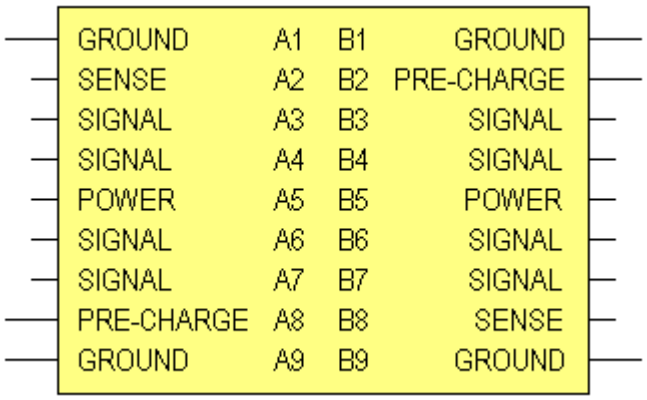


As long as the hot-swap connector is sufficiently rigid, one of the four corner pins will always be the first to engage. For a typical two-row connector arrangement this provides four first-to-make corner pins that are usually used for grounds. Other pins near the corners can be used for functions that would also benefit from this effect, for example sensing when the connector is fully seated. This diagram illustrates good practice where the grounds are in the corners and the power pins are near the center. Two sense pins are located in opposite corners so that fully seated detection is confirmed only when both of them are in contact with the slot. The remaining pins are used for all the other data signals.

## *Power electronics*

The DC power supplies to a hot-swap component are usually pre-charged by dedicated long pins that make contact before the main power pins. These pre-charge pins are protected by a circuit that limits the inrush current to an acceptable value that cannot damage the pins nor disturb the supply voltage to adjacent slots. The pre-charge circuit might be a simple series resistor, a negative temperature coefficient (NTC) resistor, or a current-limiter circuit. Further protection can be provided by a "soft-start" circuit that provides a managed ramp-up of the internal DC supply voltages within the component.

A typical sequence for a hot-swap component being plugged into a slot could be as follows:

1. Long ground pins make contact; basic electrical safety and ESD protection becomes available.
2. Long (or medium) pre-charge pins make contact; decoupling capacitors start to charge up.
3. Real time delay of tens of milliseconds.
4. Short power/signal pins make contact.
5. Connector becomes fully seated; power-on reset signal asserted within component
6. Soft-start circuit starts to apply power to the component.
7. Real time delay of tens of milliseconds.
8. Soft-start circuit completes sequence; power-on reset circuit deasserted

9. Component begins normal operation.

Hot-swap power circuits can now be purchased commercially in specially designed ASICs called hot-swap power managers (HSPMs).

## *Radio transmitters*

Modern day radio transmitters (and some TV transmitters as well) use high power RF transistor power modules instead of tubes. Hot swapping power modules is not a new technology, as many of the radio transmitters manufactured in the 1930s were capable of having power tubes swapped out while the transmitter was running—but this feature was not universally adopted due to the introduction of more reliable high power tubes.

In the mid-1990s, several radio transmitter manufactures in the US started offering swappable high power RF transistor modules.

- There was no industry standard for the design of the swappable power modules at the time.
- Early module designs had only limited patent protection.
- By the early 2000s, many transmitter models were available that used many different kinds of power modules.

The reintroduction of power modules has been good for the radio transmitter industry, as it has fostered innovation. Modular transmitters have proven to be more reliable than tube transmitters, when the transmitter is properly chosen for the conditions at the transmitting site.

Power limitations

- lowest power modular transmitter: generally 1.0 kW, using 600 W modules.
- highest power modular transmitter: 1.0 MW (for LW, MW).
- highest power modular transmitter: 45 kW (FM, TV).

Companies that produce transmitters using power modules

- Harris Broadcast (US)
- Telefunken (Germany)
- Thales (Western Europe)
- RIZ (Croatia)

## *Signal electronics*

Circuitry attached to signal pins in a hot-swap component should include some protection against electrostatic discharge (ESD). This usually takes the form of clamp diodes to ground and to the DC power supply voltage. ESD effects can be reduced by careful

design of the mechanical package around the hot-swap component, perhaps by coating it with a thin film of conductive material.

Particular care must be taken when designing systems with bussed signals which are wired to more than one hot-swap component. When a hot-swap component is inserted its input and output signal pins will represent a temporary short-circuit to ground. This can cause unwanted ground-level pulses on the signals which can disturb the operation of other hot-swap components in the system. This was a problem for early parallel SCSI disk-drives. One common design solution is to protect bussed signal pins with series diodes or resistors. CMOS buffer devices are now available with specialized inputs and outputs that minimize disturbance of bussed signals during the hot-swap operation. If all else fails, another solution is to quiesce the operation of all components during the hot-swap operation.

## *Gaming*

Hot swapping capability was more recognized among older generation video game systems. Although today many different systems can interchange games and multimedia (i.e. Blu-Ray discs) without powering down the system, this was not often the case with early-generation systems. For example, Sony PlayStation and PlayStation 2 could eject a game disc with the system powered on. However similar systems of the time, such as the Nintendo Game Boy Advance and the Nintendo 64, would freeze up and could potentially become corrupt if the game cartridge was removed with the power on; therefore manufacturers specifically warned against such practices in the owner's manual. It was supposedly the aforementioned reason Stop 'n' swop was taken out of the Banjo-Kazooie series.

## *Software*

Hot swapping can also refer to the ability to alter the running code of a program without needing to interrupt its execution. Interactive programming is a programming paradigm that makes extensive use of hot swapping, so the programming activity becomes part of the program flow itself.

Only a few programming languages support hot swapping natively, including Pike, Lisp, Erlang, Smalltalk, and Java. Microsoft Visual Studio supports a kind of hot swapping called Edit and Continue, which is supported by C#, VB.NET and C/C++ when running under a debugger.

Some web-based frameworks, such as Django, support detecting module changes and reloading them on the fly. However, although the same as hotswapping for most intents and purposes, this is technically just a cache purge, triggered by a new file. Note that this does not apply to markup and programming languages such as HTML and PHP respectively, in the general case, as these file are normally re-interpreted on each use by default. There are a few CMS's and other PHP-based frameworks (such as Drupal) that employ caching, however. In these cases, similar abilities and exceptions apply.

Hot swapping also facilitates developing systems where large amounts of data are being processed, as in entire genomes in bioinformatics algorithms.

Chapter 11

# Redundancy (Engineering) and Repetition Code

## Redundancy (engineering)

Redundant power supply

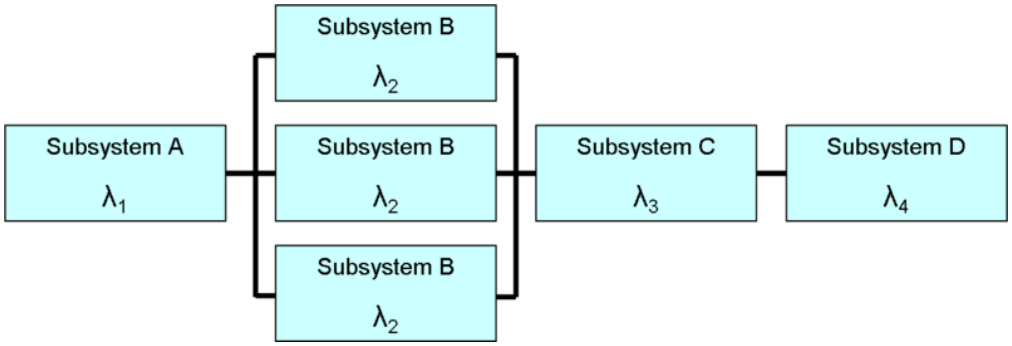


Redundant subsystem "B"

In engineering, **redundancy** is the duplication of critical components of a system with the intention of increasing reliability of the system, usually in the case of a backup or fail-safe.

In many safety-critical systems, such as fly-by-wire and hydraulic systems in aircraft, some parts of the control system may be triplicated, which is formally termed triple modular redundancy (TMR). An error in one component may then be out-voted by the other two. In a triply redundant system, the system has three sub components, all three of which must fail before the system fails. Since each one rarely fails, and the sub components are expected to fail independently, the probability of all three failing is calculated to be extremely small. Redundancy may also be known by the terms "**majority voting systems**" or "**voting logic**".

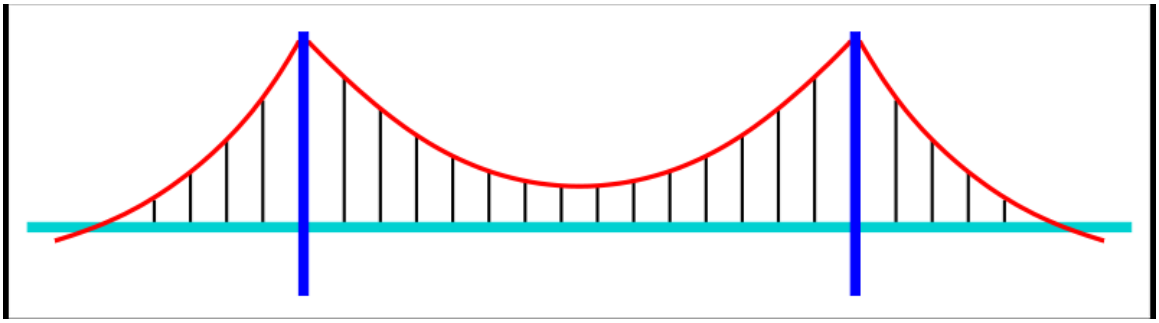

A suspension bridge's numerous cables are a form of redundancy.

## *Forms of redundancy*

There are four major forms of redundancy, these are:

- Hardware redundancy, such as DMR and TMR
- Information redundancy, such as Error detection and correction methods
- Time redundancy, including transient fault detection methods such as **Alternate Logic**
- Software redundancy such as N-version programming

## *Function of redundancy*

The two functions of redundancy are passive redundancy and active redundancy. Both functions prevent performance decline from exceeding specification limits without human intervention using extra capacity.

Passive redundancy uses excess capacity to reduce the impact of component failures. One common form of passive redundancy is the extra strength of cabling and struts used in bridges. This extra strength allows some structural components to fail without bridge collapse. The extra strength used in the design is called the margin of safety.

Eyes and ears provide working examples of passive redundancy. Vision loss in one eye does not cause blindness but depth perception is impaired. Hearing loss in one ear does not cause deafness but directionality is impaired. Performance decline is commonly associated with passive redundancy when a limited number of failures occur.

Active redundancy eliminates performance decline by monitoring performance of individual device, and this monitoring is used in voting logic. The voting logic is linked to switching that automatically reconfigures components. Error detection and correction and the Global Positioning System (GPS) are two examples of active redundancy.

Electrical power distribution provides an example of active redundancy. Several power lines connect each generation facility with customers. Each power line include monitors that detect overload. Each power line also includes circuit breakers. The combination of power lines provides excess capacity. Circuit breakers disconnect a power line when monitors detect an overload. Power is redistributed across the remaining lines.

## *Voting Logic*

Voting logic uses performance monitoring to determine how to reconfigure individual components so that operation continues without violating specification limitations of the overall system. Voting logic often involve computers, but systems composed of items other than computers may be reconfigured using voting logic. Circuit breakers are an example of a form of non-computer voting logic.

Electrical power systems use power scheduling to reconfigure active redundancy. Computing systems adjust the production output of each generating facility when other generating facilities are suddenly lost. This prevents blackout conditions during major events like earthquake.

The simplest voting logic in computing systems involves two components: primary and alternate. They both run similar software, but the output from the alternate remains inactive during normal operation. The primary monitors itself and periodically sends an activity message to the alternate as long as everything is OK. All outputs from the primary stop, including the activity message, when the primary detects a fault. The alternate activates its output and takes over from the primary after a brief delay when the

activity message ceases. Errors in voting logic can cause both to have all outputs active at the same time, can cause both to have all outputs inactive at the same time, or outputs can flutter on and off.

A more reliable form of voting logic involves an odd number of 3 devices or more. All perform identical functions and the outputs are compared by the voting logic. The voting logic establishes a majority when there is a disagreement, and the majority will act to deactivate the output from other device(s) that disagree. A single fault will not interrupt normal operation. This technique is used with avionics systems, such as those responsible for operation of the space shuttle.

## *Calculating the probability of system failure*

Each duplicate component added to the system decreases the probability of system failure according to the formula:

$$p = \prod_{i=1}^{n} p_i$$

where:

- $n$ - number of components
- $p_i$ - probability of component i failing
- $p$ - the probability of all components failing (system failure)

This formula assumes independence of failure events. That means that the probability of a component B failing given that a component A has already failed is the same as that of B failing when A has not failed. There are situations where this is unreasonable, such as using two power supplies connected to the same socket, whereby if one socket failed, the other would too.

It also assumes that at only one component is needed to keep the system running. If $m$ components are needed for the system to survive, out of $n$, the probability of failure is

$P = 1 - ((1-p)^{(n-m)} C_n^m)$, Assuming all components have equal probability, $p$, of failure

This model is probably unrealistic in that it assumes that components are not replaced in time when they fail.

# Repetition code

**Repetition code** is an (*r*,1) coding scheme that repeats the bits across a channel to achieve error free communication (r is the number of bits in each codeword for each data bit to be coded). Repetition code is generally a very naive method of encoding data across a channel, and it is not preferred for *Additive White Gaussian Noise Channels* (AWGN), due to its worse-than-the-present error performance. Repetition codes generally offer a poor compromise between data rate and bit error rate, and other forms of error correcting codes can provide superior performance in these areas. The chief attraction of the repetition code is the ease of implementation.

There are two parts to the repetition code, as for any other code: the encoder and decoder, which will be described in detail.

## *Repetition Coder*

The encoder is a simple device that repeats, *r* times, a particular bit to the waveform modulator when the bit is received from the source stream.

For example, if we have a (3,1) repetition code, then encoding the signal $m = 101001$ yields a code $c = 111000111000000111$.

## *Repetition Decoder*

Repetition decoding is usually done using Majority logic detection. To determine the value of a particular bit, we look at the received copies of the bit in the stream and choose the value that occurs more frequently.

For example, suppose we have a (3,1) repetition code and we are decoding the signal $c = 110001111$. The decoded message is $m = 101$, as we have most occurrence of 1's (two to one), 0's (two to one), and 1's (three to zero) in the first, second, and third code sequences, respectively.

This approach discards any 'soft' probability information obtained when decoding each received bit, and the performance of the code can be improved by retaining this probability information and using it to derive a joint probability across all n bits of the actual information bit value.

## *Repetition Codes on Fading Channel*

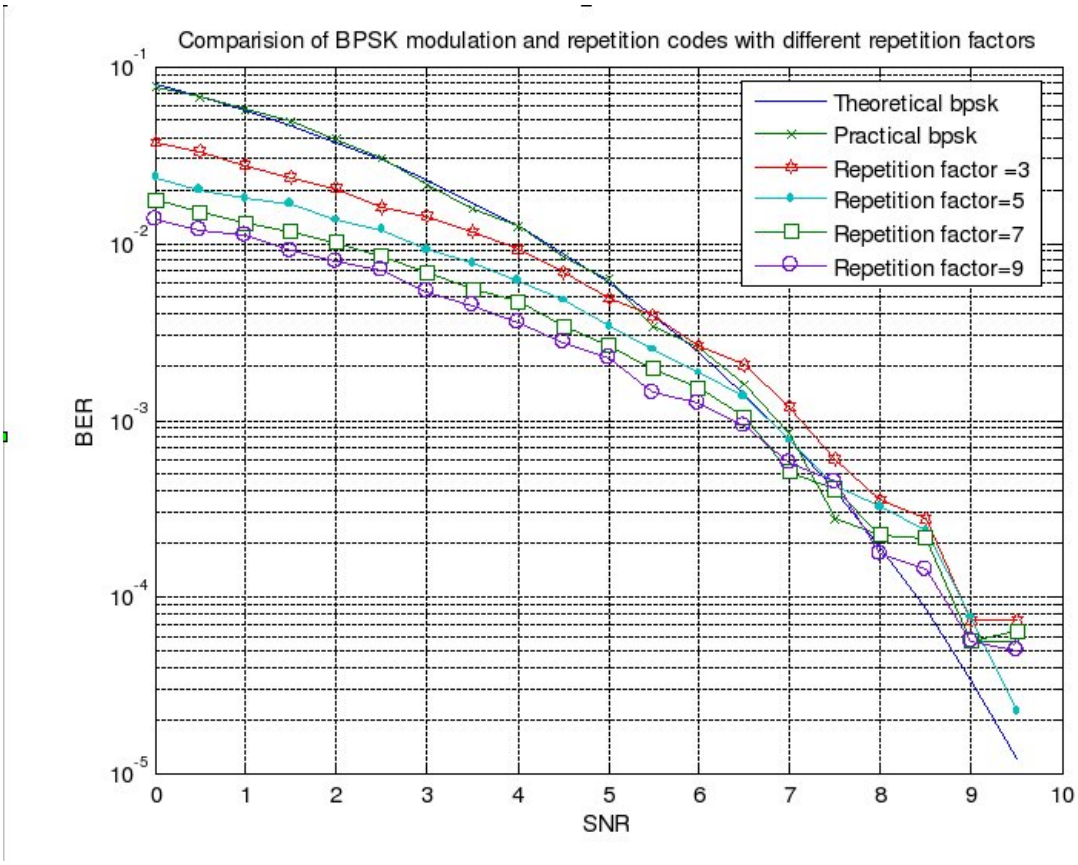


For fading channels repetition codes perform well with increasing repetition factor. In this figure, the coding gains for various repetition factors are seen.

## *Repetition Codes on Gaussian Channel*

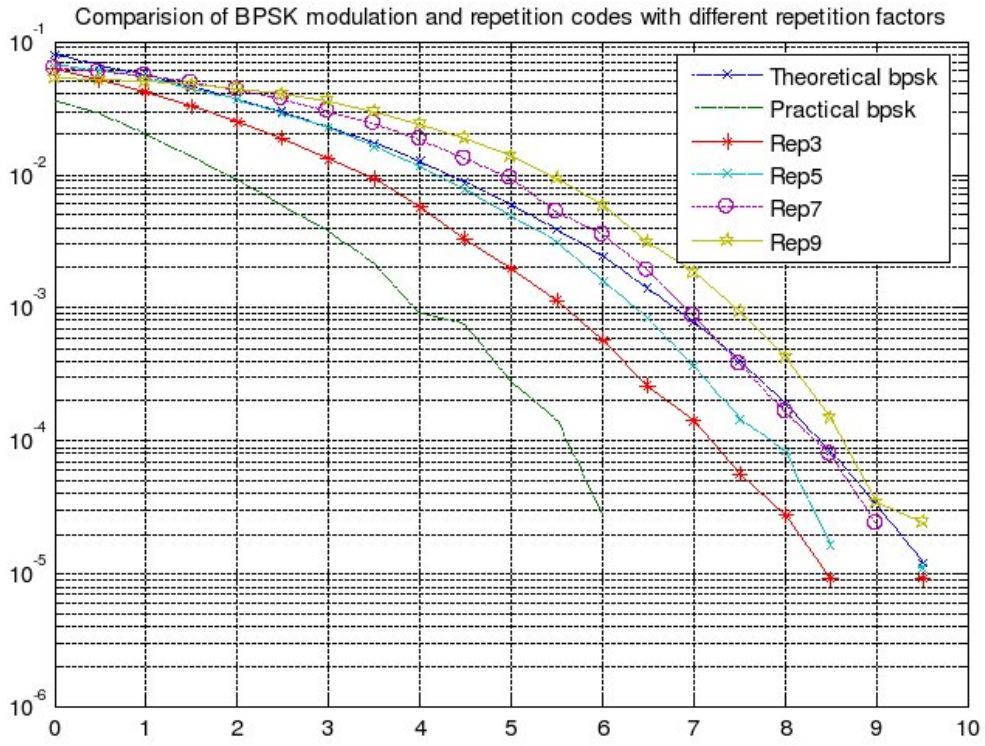


For the AWGN channels perform worse for longer repetition factors. In this figure, the coding gains are progressively worse with the increasing parameter.

## *Code parameters*

The minimum Hamming distance ($d_{min}$) is $r$ for an ($r$,1) repetition code, and there are two valid code words - all ones and all zeros, so the minimum weight ($w_{min}$) is r. This gives the repetition code an error correcting capacity of $\frac{r}{2}$(i.e. it will correct up to $\frac{r}{2}$errors in any code word).

## *Applications*

Due to the simplicity of the channel encoding and decoding for repetition codes, they find applications in fading channels and non-AWGN environments. Repetition codes can be viewed as a method of space-time diversity as well.

Most modulation techniques transmit a bit or chip over many cycles of a sinusoid carrier signal. The low-pass filter used to average the relevant parameter (amplitude, phase, or frequency) over the entire bit-time or chip-time can be seen as a kind of repetition decoder.

Some UARTs, such as the ones used in the FlexRay protocol, use a majority filter to ignore brief noise spikes. This spike-rejection filter can be seen as a kind of repetition decoder.

Despite their poor performance as stand-alone codes, use in Turbo code-like iteratively decoded concatenated coding schemes, such as repeat-accumulate (RA) and accumulate-repeat-accumulate (ARA) codes, allows for surprisingly good error correction performance.

Repetition codes are one of the few known codes whose code rate can be automatically adjusted to varying channel capacity, by sending more or less parity information as required to overcome the channel noise, and it is the only such code known for non-erasure channels. Practical adaptive codes for erasure channels have been invented only recently, and are known as fountain codes.

# Chapter 12

# Backup

In Information Technology, a **backup** or the process of **backing up** refers to making copies of data so that these additional copies may be used to *restore* the original after a data loss event. The verb form is *back up* in two words, whereas the noun is *backup* (often used like an adjective in compound nouns).

Backups have two distinct purposes. The primary purpose is to recover data as a reaction to data loss, be it by data deletion or corrupted data. Data loss is a very common experience of computer users. 67% of internet users have suffered serious data loss. The secondary purpose of backups is to recover data from a historical period of time within the constraints of a user-defined data retention policy, typically configured within a backup application for how long copies of data are required. Though backups popularly represent a simple form of disaster recovery, and should be part of a disaster recovery plan, by themselves, backups should not alone be considered disaster recovery. Not all backup systems and/or backup applications are able to reconstitute a computer system, or in turn other complex configurations such as a computer cluster, active directory servers, or a database server, by restoring only data from a backup.

Since a backup system contains at least one copy of all data worth saving, the data storage requirements are considerable. Organizing this storage space and managing the backup process is a complicated undertaking. A data repository model can be used to provide structure to the storage. In the modern era of computing there are many different types of data storage devices that are useful for making backups. There are also many different ways in which these devices can be arranged to provide geographic redundancy, data security, and portability.

Before data is sent to its storage location, it is selected, extracted, and manipulated. Many different techniques have been developed to optimize the backup procedure. These include optimizations for dealing with open files and live data sources as well as compression, encryption, and de-duplication, among others. Many organizations and individuals try to have confidence that the process is working as expected and work to define measurements and validation techniques. It is also important to recognize the limitations and human factors involved in any backup scheme.

## *Storage, the base of a backup system*

### Data repository models

Any backup strategy starts with a concept of a data repository. The backup data needs to be stored somehow and probably should be organized to a degree. It can be as simple as a sheet of paper with a list of all backup tapes and the dates they were written or a more sophisticated setup with a computerized index, catalog, or relational database. Different repository models have different advantages. This is closely related to choosing a backup rotation scheme.

**Unstructured**

An unstructured repository may simply be a stack of floppy disks or CD-R/DVD-R media with minimal information about what was backed up and when. This is the easiest to implement, but probably the least likely to achieve a high level of recoverability.

**Full + incrementals**

A full + incremental repository aims to make it more feasible to store several copies of the source data. At first, a *full* backup (of all files) is made. After that, any number of *incremental* backups can be made. There are many different types of incremental backups, but they all attempt to only back up a small amount of data (when compared to the size of a full backup). An incremental backup copies everything that has changed since the last backup (full, differential or incremental). Restoring a whole system to a certain point in time would require locating the last full backup taken previous to that time and all the incremental backups that cover the period of time between the full backup and the particular point in time to which the system is supposed to be restored. The scope of an incremental backup is typically defined as the period of time between other full or incremental backups. Different implementations of backup systems frequently use specialized or conflicting definitions of these terms.

**Differential backup**

A differential backup copies files that have been created or changed since the last full backup. It does not mark files as having been backed up (in other words, the archive attribute is not cleared). If you are performing a combination of full and differential backups, restoring files and folders requires that you have the last full as well as the last differential backup.

**Reverse delta**

A reverse delta system stores the differences between current versions of a system and previous versions. A reverse delta backup will start with a normal full backup. After the full backup is performed, the system will periodically synchronize the full backup with the live copy, while storing the data necessary to reconstruct older versions. This can either be done using hard links, or using binary diffs. This system works particularly well for large, slowly changing, data sets. Examples of programs that use this method are rdiff-backup and Time Machine

**Continuous data protection**

Instead of scheduling periodic backups, the system immediately logs every change on the host system. This is generally done by saving byte or block-level differences rather than file-level differences. It differs from simple disk mirroring in that it enables a roll-back of the log and thus restoration of old image of data.

**Full system backup**

This type of backup is designed to allow an entire PC to be recovered to "bare metal" without any installation of operating system, application software and data. Most users understand that a backup will prevent "data" from being lost. The expense in a full system recovery is in the hours that it takes for a technician to rebuild a machine to the point of restoring the last data backup. So, a full system backup makes a complete image of the computer so that if needed, it can be copied back to the PC, usually using some type of bespoke software such as Ghost, and the user can carry on from that point.

## Storage media

Regardless of the repository model that is used, the data has to be stored on some data storage medium somewhere.

Magnetic tape

Magnetic tape has long been the most commonly used medium for bulk data storage, backup, archiving, and interchange. Tape has typically had an order of magnitude better capacity/price ratio when compared to hard disk, but recently the ratios for tape and hard disk have become a lot closer. There are myriad formats, many of which are proprietary or specific to certain markets like mainframes or a particular brand of personal computer. Tape is a sequential access medium, so even though access times may be poor, the rate of continuously writing or reading data can actually be very fast. Some new tape drives are even faster than modern hard disks. A principal advantage of tape is that it has been used for this purpose for decades (much longer than any alternative) and its characteristics are well understood.

Hard disk

The capacity/price ratio of hard disk has been rapidly improving for many years. This is making it more competitive with magnetic tape as a bulk storage medium. The main advantages of hard disk storage are low access times, availability, capacity and ease of use. External disks can be connected via local interfaces like SCSI, USB, FireWire, or eSATA, or via longer distance technologies like Ethernet, iSCSI, or Fibre Channel. Some disk-based backup systems, such as Virtual Tape Libraries, support data deduplication which can dramatically reduce the amount of disk storage capacity consumed by daily and weekly backup data. The main disadvantages of hard disk backups are that they are easily damaged, especially while being transported (e.g., for off-site backups), and that their stability over periods of years is a relative unknown.

Optical storage

Blu-ray Discs dramatically increase the amount of data possible on a single optical storage disk. Systems containing Blu-ray discs can store massive amounts of data and be more cost efficient than hard drives and magnetic tape. Some optical storage systems allow for cataloged data backups without human contact with the discs, allowing for longer data integrity. A recordable CD can be used as a backup device. One advantage of CDs is that they can be restored on any machine with a CD-ROM drive. (In practice, writable CD-ROMs are not always universally readable.) In addition, recordable CD's are relatively cheap. Another common format is recordable DVD. Many optical disk formats are WORM type, which makes them useful for archival purposes since the data can't be changed. Other rewritable formats can also be utilized such as CD-RW or DVD-RAM.

Floppy disk

During the 1980s and early 1990s, many personal/home computer users associated backing up mostly with copying to floppy disks. The low data capacity of a floppy disk makes it an unpopular and obsolete choice today.

Solid state storage

Also known as flash memory, thumb drives, USB flash drives, CompactFlash, SmartMedia, Memory Stick, Secure Digital cards, etc., these devices are relatively costly for their low capacity, but offer excellent portability and ease-of-use.

Remote backup service

As broadband internet access becomes more widespread, remote backup services are gaining in popularity. Backing up via the internet to a remote location can protect against some worst-case scenarios such as fires, floods, or earthquakes which would destroy any backups in the immediate vicinity along with everything else. There are, however, a number of drawbacks to remote backup services. First, Internet connections are usually slower than local data storage devices. Residential broadband is especially problematic as routine backups must use an upstream link that's usually much slower than the downstream link used only occasionally to retrieve a file from backup. This tends to limit the use of such services to relatively small amounts of high value data. Secondly, users must trust a third party service provider to maintain the privacy and integrity of their data, although confidentiality can be assured by encrypting the data before transmission to the backup service with an encryption key known only to the user. Ultimately the backup service must itself use one of the above methods so this could be seen as a more complex way of doing traditional backups.

## Managing the data repository

Regardless of the data repository model or data storage media used for backups, a balance needs to be struck between accessibility, security and cost. These media management methods are not mutually exclusive and are frequently combined to meet the needs of the situation. Using on-line disks for staging data before it is sent to a near-line tape library is a common example.

On-line

On-line backup storage is typically the most accessible type of data storage, which can begin restore in milliseconds time. A good example would be an internal hard disk or a disk array (maybe connected to SAN). This type of storage is very convenient and speedy, but is relatively expensive. On-line storage is quite vulnerable to being deleted or overwritten, either by accident, by intentional malevolent action, or in the wake of a data-deleting virus payload.

Near-line
: Near-line storage is typically less accessible and less expensive than on-line storage, but still useful for backup data storage. A good example would be a tape library with restore times ranging from seconds to a few minutes. A mechanical device is usually involved in moving media units from storage into a drive where the data can be read or written. Generally it has safety properties similar to on-line storage.

Off-line
: Off-line storage requires some direct human action in order to make access to the storage media physically possible. This action is typically inserting a tape into a tape drive or plugging in a cable that allows a device to be accessed. Because the data is not accessible via any computer except during limited periods in which it is written or read back, it is largely immune to a whole class of on-line backup failure modes. Access time will vary depending on whether the media is on-site or off-site.

Off-site data protection
: To protect against a disaster or other site-specific problem, many people choose to send backup media to an off-site vault. The vault can be as simple as a system administrator's home office or as sophisticated as a disaster-hardened, temperature-controlled, high-security bunker that has facilities for backup media storage. Importantly a data replica *can* be off-site but also *on-line* (e.g., an off-site RAID mirror). Such a replica has fairly limited value as a backup, and should not be confused with an off-line backup.

Backup site or disaster recovery center (DR center)
: In the event of a disaster, the data on backup media will not be sufficient to recover. Computer systems onto which the data can be restored and properly configured networks are necessary too. Some organizations have their own data recovery centers that are equipped for this scenario. Other organizations contract this out to a third-party recovery center. Because a DR site is itself a huge investment, backing up is very rarely considered the preferred method of moving data to a DR site. A more typical way would be remote disk mirroring, which keeps the DR data as up to date as possible.

## *Selection and extraction of data*

A successful backup job starts with selecting and extracting coherent units of data. Most data on modern computer systems is stored in discrete units, known as files. These files are organized into filesystems. Files that are actively being updated can be thought of as "live" and present a challenge to back up. It is also useful to save metadata that describes the computer or the filesystem being backed up.

Deciding what to back up at any given time is a harder process than it seems. By backing up too much redundant data, the data repository will fill up too quickly. Backing up an insufficient amount of data can eventually lead to the loss of critical information.

## Files

Copying files
: Making copies of files is the simplest and most common way to perform a backup. A means to perform this basic function is included in all backup software and all operating systems.

Partial file copying
: Instead of copying whole files, one can limit the backup to only the blocks or bytes within a file that have changed in a given period of time. This technique can use substantially less storage space on the backup medium, but requires a high level of sophistication to reconstruct files in a restore situation. Some implementations require integration with the source filesystem.

When backing up over a network, the rsync utility automatically transmits a minimum set of changes to bring an earlier version of a file at the destination up to date with the current version at the source. Rsync can dramatically reduce the network traffic needed to maintain a remote mirror of a large set of files undergoing small, frequent changes.

## Filesystems

Filesystem dump
: Instead of copying files within a filesystem, a copy of the whole filesystem itself can be made. This is also known as a *raw partition backup* and is related to disk imaging. The process usually involves unmounting the filesystem and running a program like dd (Unix). Because the disk is read sequentially and with large buffers, this type of backup can be much faster than reading every file normally, especially when the filesystem contains many small files, is highly fragmented, or is nearly full. But because this method also reads the free disk blocks that contain no useful data, this method can also be slower than conventional reading, especially when the filesystem is nearly empty. Some filesystems, such as XFS, provide a "dump" utility that reads the disk sequentially for high performance while skipping unused sections. The corresponding restore utility can selectively restore individual files or the entire volume at the operator's choice.

Identification of changes
: Some filesystems have an archive bit for each file that says it was recently changed. Some backup software looks at the date of the file and compares it with the last backup to determine whether the file was changed.

Versioning file system
: A versioning filesystem keeps track of all changes to a file and makes those changes accessible to the user. Generally this gives access to any previous version, all the way back to the file's creation time. An example of this is the Wayback versioning filesystem for Linux.

If a computer system is in use while it is being backed up, the possibility of files being open for reading or writing is real. If a file is open, the contents on disk may not correctly represent what the owner of the file intends. This is especially true for database files of all kinds. The term fuzzy backup can be used to describe a backup of live data that looks like it ran correctly, but does not represent the state of the data at any single point in time. This is because the data being backed up changed in the period of time between when the backup started and when it finished. For databases in particular, fuzzy backups are worthless.

Snapshot backup
: A snapshot is an instantaneous function of some storage systems that presents a copy of the file system as if it were frozen at a specific point in time, often by a copy-on-write mechanism. An effective way to back up live data is to temporarily quiesce it (e.g. close all files), take a snapshot, and then resume live operations. At this point the snapshot can be backed up through normal methods. While a snapshot is very handy for viewing a filesystem as it was at a different point in time, it is hardly an effective backup mechanism by itself.

Open file backup
: Many backup software packages feature the ability to handle open files in backup operations. Some simply check for openness and try again later. File locking is useful for regulating access to open files.
: When attempting to understand the logistics of backing up open files, one must consider that the backup process could take several minutes to back up a large file such as a database. In order to back up a file that is in use, it is vital that the entire backup represent a single-moment snapshot of the file, rather than a simple copy of a read-through. This represents a challenge when backing up a file that is constantly changing. Either the database file must be locked to prevent changes, or a method must be implemented to ensure that the original snapshot is preserved long enough to be copied, all while changes are being preserved. Backing up a file while it is being changed, in a manner that causes the first part of the backup to represent data *before* changes occur to be combined with later parts of the backup *after* the change results in a corrupted file that is unusable, as most large files contain internal references between their various parts that must remain consistent throughout the file.

Cold database backup
: During a cold backup, the database is closed or locked and not available to users. The datafiles do not change during the backup process so the database is in a consistent state when it is returned to normal operation.

Hot database backup
: Some database management systems offer a means to generate a backup image of the database while it is online and usable ("hot"). This usually includes an inconsistent image of the data files plus a log of changes made while the procedure is running. Upon a restore, the changes in the log files are reapplied to bring the database in sync.

### Metadata

Not all information stored on the computer is stored in files. Accurately recovering a complete system from scratch requires keeping track of this non-file data too.

System description
: System specifications are needed to procure an exact replacement after a disaster.

Boot sector
: The boot sector can sometimes be recreated more easily than saving it. Still, it usually isn't a normal file and the system won't boot without it.

Partition layout
: The layout of the original disk, as well as partition tables and filesystem settings, is needed to properly recreate the original system.

File metadata
: Each file's permissions, owner, group, ACLs, and any other metadata need to be backed up for a restore to properly recreate the original environment.

System metadata
: Different operating systems have different ways of storing configuration information. Microsoft Windows keeps a registry of system information that is more difficult to restore than a typical file.

## *Manipulation of data and dataset optimization*

It is frequently useful or required to manipulate the data being backed up to optimize the backup process. These manipulations can provide many benefits including improved backup speed, restore speed, data security, media usage and/or reduced bandwidth requirements.

Compression
: Various schemes can be employed to shrink the size of the source data to be stored so that it uses less storage space. Compression is frequently a built-in feature of tape drive hardware.

De-duplication
: When multiple similar systems are backed up to the same destination storage device, there exists the potential for much redundancy within the backed up data. For example, if 20 Windows workstations were backed up to the same data repository, they might share a common set of system files. The data repository only needs to store one copy of those files to be able to restore any one of those workstations. This technique can be applied at the file level or even on raw blocks of data, potentially resulting in a massive reduction in required storage space. Deduplication can occur on a server before any data moves to backup media, sometimes referred to as source/client side deduplication. This approach also reduces bandwidth required to send backup data to its target media. The process can also occur at the target storage device, sometimes referred to as inline or back-end deduplication.

Duplication

Sometimes backup jobs are duplicated to a second set of storage media. This can be done to rearrange the backup images to optimize restore speed or to have a second copy at a different location or on a different storage medium.

Encryption
: High capacity removable storage media such as backup tapes present a data security risk if they are lost or stolen. Encrypting the data on these media can mitigate this problem, but presents new problems. Encryption is a CPU intensive process that can slow down backup speeds, and the security of the encrypted backups is only as effective as the security of the key management policy.

Multiplexing
: When there are many more computers to be backed up than there are destination storage devices, the ability to use a single storage device with several simultaneous backups can be useful.

Refactoring
: The process of rearranging the backup sets in a data repository is known as refactoring. For example, if a backup system uses a single tape each day to store the incremental backups for all the protected computers, restoring one of the computers could potentially require many tapes. Refactoring could be used to consolidate all the backups for a single computer onto a single tape. This is especially useful for backup systems that do *incrementals forever* style backups.

Staging
: Sometimes backup jobs are copied to a staging disk before being copied to tape. This process is sometimes referred to as D2D2T, an acronym for Disk to Disk to Tape. This can be useful if there is a problem matching the speed of the final destination device with the source device as is frequently faced in network-based backup systems. It can also serve as a centralized location for applying other data manipulation techniques.

## *Managing the backup process*

It is important to understand that backing up is a process. As long as new data is being created and changes are being made, backups will need to be updated. Individuals and organizations with anything from one computer to thousands (or even millions) of computer systems all have requirements for protecting data. While the scale is different, the objectives and limitations are essentially the same. Likewise, those who perform backups need to know to what extent they were successful, regardless of scale.

### Objectives

Recovery point objective (RPO)
: The point in time that the restarted infrastructure will reflect. Essentially, this is the roll-back that will be experienced as a result of the recovery. The most desirable RPO would be the point just prior to the data loss event. Making a more recent recovery point achievable requires increasing the frequency of synchronization between the source data and the backup repository.

Recovery time objective (RTO)

The amount of time elapsed between disaster and restoration of business functions.

Data security
: In addition to preserving access to data for its owners, data must be restricted from unauthorized access. Backups must be performed in a manner that does not compromise the original owner's undertaking. This can be achieved with data encryption and proper media handling policies.

## Limitations

An effective backup scheme will take into consideration the limitations of the situation.

Backup window
: The period of time when backups are permitted to run on a system is called the backup window. This is typically the time when the system sees the least usage and the backup process will have the least amount of interference with normal operations. The backup window is usually planned with users' convenience in mind. If a backup extends past the defined backup window, a decision is made whether it is more beneficial to abort the backup or to lengthen the backup window.

Performance impact
: All backup schemes have some performance impact on the system being backed up. For example, for the period of time that a computer system is being backed up, the hard drive is busy reading files for the purpose of backing up, and its full bandwidth is no longer available for other tasks. Such impacts should be analyzed.

Costs of hardware, software, labor
: All types of storage media have a finite capacity with a real cost. Matching the correct amount of storage capacity (over time) with the backup needs is an important part of the design of a backup scheme. Any backup scheme has some labor requirement, but complicated schemes have considerably higher labor requirements. The cost of commercial backup software can also be considerable.

Network bandwidth
: Distributed backup systems can be affected by limited network bandwidth.

## Implementation

Meeting the defined objectives in the face of the above limitations can be a difficult task. The tools and concepts below can make that task more achievable.

Scheduling
: Using a job scheduler can greatly improve the reliability and consistency of backups by removing part of the human element. Many backup software packages include this functionality.

Authentication

Over the course of regular operations, the user accounts and/or system agents that perform the backups need to be authenticated at some level. The power to copy all data off of or onto a system requires unrestricted access. Using an authentication mechanism is a good way to prevent the backup scheme from being used for unauthorized activity.

Chain of trust
: Removable storage media are physical items and must only be handled by trusted individuals. Establishing a chain of trusted individuals (and vendors) is critical to defining the security of the data.

## Measuring the process

To ensure that the backup scheme is working as expected, the process needs to include monitoring key factors and maintaining historical data.

Backup validation
: (also known as "backup success validation") The process by which owners of data can get information about how their data was backed up. This same process is also used to prove compliance to regulatory bodies outside of the organization, for example, an insurance company might be required under HIPAA to show "proof" that their patient data are meeting records retention requirements. Disaster, data complexity, data value and increasing dependence upon ever-growing volumes of data all contribute to the anxiety around and dependence upon successful backups to ensure business continuity. For that reason, many organizations rely on third-party or "independent" solutions to test, validate, and optimize their backup operations (backup reporting).

Reporting
: In larger configurations, reports are useful for monitoring media usage, device status, errors, vault coordination and other information about the backup process.

Logging
: In addition to the history of computer generated reports, activity and change logs are useful for monitoring backup system events.

Validation
: Many backup programs make use of checksums or hashes to validate that the data was accurately copied. These offer several advantages. First, they allow data integrity to be verified without reference to the original file: if the file as stored on the backup medium has the same checksum as the saved value, then it is very probably correct. Second, some backup programs can use checksums to avoid making redundant copies of files, to improve backup speed. This is particularly useful for the de-duplication process.

Monitored backup
: Backup processes are monitored by a third party monitoring center. This center alerts users to any errors that occur during automated backups. Monitored backup requires software capable of pinging the monitoring center's servers in the case of errors. Some monitoring services also allow collection of historical meta-data, that can be used for Storage Resource Management purposes like projection of

data growth, locating redundant primary storage capacity and reclaimable backup capacity. The Wizards Storage Portal is an example of a solution that monitors IBM's well known Tivoli Storage Manager(TSM) solution.

## *Lore*

### Confusion

Due to a considerable overlap in technology, backups and backup systems are frequently confused with archives and fault-tolerant systems. Backups differ from archives in the sense that archives are the *primary copy* of data, usually put away for future use, while backups are a *secondary copy* of data, kept on hand to replace the original item. Backup systems differ from fault-tolerant systems in the sense that backup systems assume that a fault *will* cause a data loss event and fault-tolerant systems assure a fault *will not*.

### Advice

- The more important the data that is stored on the computer, the greater is the need for backing up this data.
- A backup is only as useful as its associated restore strategy. For critical systems and data, the restoration process must be tested.
- Storing the copy near the original is unwise, since many disasters such as fire, flood, theft, and electrical surges are likely to cause damage to the backup at the same time. In these cases, both the original and the backup medium are likely to be lost.
- Automated backup and scheduling should be considered, as manual backups can be affected by human error.
- Backups can fail for a wide variety of reasons. A verification or monitoring strategy is an important part of a successful backup plan.
- Multiple backups on different media, stored in different locations, should be used for all critical information.
- Backed up archives should be stored in open and standard formats, especially when the goal is long-term archiving. Recovery software and processes may have changed, and software may not be available to restore data saved in proprietary formats.
- System administrators and others working in the information technology field are routinely fired for not devising and maintaining backup processes suitable to their organization.
- If you already have a tape backup system, a second backup program may be necessary. Perform an additional backup to an external hard disk with an automatic backup program so you will have doubled the data security, and it is easy to check the backed-up files in the external hard disk.

## Events

- In 1996, during a fire at the headquarters of Crédit Lyonnais, a major bank in Paris, system administrators ran into the burning building to rescue backup tapes because they didn't have off-site copies. Crucial bank archives and computer data were lost.
- Privacy Rights Clearinghouse has documented 16 instances of stolen or lost backup tapes (among major organizations) in 2005 & 2006. Affected organizations included Bank of America, Ameritrade, Citigroup, and Time Warner.
- On 3 January 2008, an email server crashed at TeliaSonera, a major Nordic telecom company and internet service provider. It was subsequently discovered that the last serviceable backup set was from 15 December 2007. Three hundred thousand customer email accounts were affected.
- On 27 February 2011 a software bug on Gmail caused 0.02% of its users to lose all their emails. The messages were successfully restored from tape backups hours after the event.