



Željko Nastasić
Gábor Deák Jahn

The Citroën Technical Guide

There are many car manufacturers, makes, models and versions on the road today but—as we all know—none of them compares to Citroën in its engineering excellence, especially regarding suspension comfort, roadholding, and stability.

In this book we tried to describe how the various subsystems work. We never intended to replace service manuals or similar technical instructions. Illustrations are schematic, focusing on the principles of operation rather than on minute details of implementation.

This guide is not linked to any specific Citroën model but describes all systems and solutions used on a large number of cars from the glorious line of DS, ID, CX, GS, GSA, BX, XM, Xantia, Xsara and the C5.

Table of Contents

Fuel Injection

Electronic Fuel Injection	6
The HPi engine	11
Diesel engines	12
Electronic Diesel Control	17
Diesel Direct Injection	19

Suspension

A Suspension Primer	22
Hydropneumatic Suspension	23
Hydractive I	27
Hydractive II	32
Anti-sink system	35
Activa Suspension	36
Hydractive 3	38

Steering

Power Assisted Steering	42
-----------------------------------	----

DIRAVI Steering	44
Self-steering Rear	46

Brakes

Standard braking system	48
Anti-lock Braking System	50

Electrical Systems

Multiplex network	52
-----------------------------	----

Air Conditioning

Air conditioning	56
----------------------------	----

Appendix


ORGA number	60
-----------------------	----

Index

We would like to thank the following people for their valuable contribution:

✉ Adam Reif (HPi diesel engine)

Fuel Injection



Electronic Fuel Injection

The Otto engine needs a mixture of fuel and air for its operation. It would be the task of the fuel supply—carburetor or injection—to provide the engine with the ideal mixture. Unfortunately, there is no such thing as an ideal mixture.

Perfect combustion, as chemistry calls it, would require air and fuel in proportion of 14.7 parts to 1 (this is the so-called stoichiometric ratio). While this might be satisfactory for the scientists, the real-life conditions of a vehicle call for slightly different characteristics.

We use the ratio of actual mixture to the stoichiometric mixture, called *lambda* (λ), to describe the composition of the mixture entering the engine: $\lambda=1$ denotes the chemically ideal mixture, $\lambda<1$ means rich, $\lambda>1$ is lean.

The best performance would require a slightly rich mixture, with the lambda around 0.9, while fuel economy would call for a slightly lean one, between 1.1 and 1.3. Some harmful components in exhaust gas would reduce in quantity between lambda values of 1 to 1.2, others below 0.8 or above 1.4. And if this is not yet enough, a cold engine requires a very rich mixture to keep running. After

warming up, the mixture can return to normal, but the temperature of the incoming air still plays a significant role: the cooler the air, the denser it becomes, and this influences the lambda ratio as well.

All these requirements are impossible to satisfy with simpler mechanical devices like carburetors. Electronic fuel injection provides a system that can measure the many circumstances the engine is operating in and decide on the amount of fuel (in other words, the lambda ratio) entering the engine. By carefully adjusting the internal rules of this device, manufacturers can adapt the characteristics of the fuel injection to the actual requirements: a sporty GTi would demand rather different settings than a city car; besides, catalytic converters have their own demands that, as we will later see, upset the appletart quite vehemently.

Earlier fuel injection systems only knew about fuel, the ignition was supplied by traditional methods. Later on, these systems (now called engine management systems) took on the duty of generating the sparks as well. But even with this second incarnation, the fuel injection part remained practically the same, thus the following section applies to both kind of systems.

Fuel injection

The two most important inputs describing the actual operating condition of the engine, thus determining the fuel demand are the **engine speed** (revolution) and **engine load**. The engine speed can be measured easily on systems using traditional ignition: the ignition primary circuit generates pulses with their frequency proportional to engine speed (the tachometer uses this same signal to show the rpm to the driver). When the injection system provides the ignition as well, it cannot at the same time rely on it, so an additional sensor is used instead.

The engine load is usually determined by measuring the quantity of air the engine tries to suck in. There are various methods of attaining this: earlier systems used a flap which is deflected by the air flowing through the sensor—the angle of deflection is proportional to the amount of air passing through (**air flow sensor, AFS**). Later systems used a pressure sensor measuring the pressure inside the inlet manifold (**manifold absolute pressure, MAP sensor**). Yet another system (although not used on Citroëns) heats a platinum wire and lets the incoming air passing around cool it; by measuring the current needed to keep the wire temperature at a constant value above the temperature of the incoming air, the mass of air can be determined. Some simpler systems do not even measure the amount of air but use a pre-stored table in their computer to approximate it based upon the engine speed and the position of the throttle pedal—not that accurate but certainly much cheaper.

Under ideal conditions, these two inputs would already be enough to control the engine. A large table can be set


Amount of fuel injected		Engine load			
		0%	5%	...	100%
Engine speed	idle	3	3	...	3
	850 rpm	4	5	...	5
	900 rpm	5	6	...	7
	...	...	...	...	...
	6,000 rpm	9	8	...	10

up, like the one illustrated here (of course, this is an illustration only, the actual values mean nothing here), and for any pair of incoming engine speed and load values the necessary fuel amount can be determined. By keeping the pressure of fuel constant behind the injector valves, the amount of fuel injected depends solely on the time period the injectors are opened for, hence, the table can contain injector opening times.

An this is exactly how it is done in modern injection systems: the controlling microcomputer keeps a lookup table like this to determine the base pulse width. Earlier systems were constructed from discrete, analog elements, not like a small computer; a more or less equivalent circuit made of various hybrid resistance arrays and semiconductors were used for the same purpose.

Chip tuning, by the way, is the simple operation of replacing the said table with another one, yielding different characteristics (usually to gain power, allowing for worse fuel economy). As the computer stores this table in a programmable memory—similar in function to the BIOS in desktop computers—, replacing it is possible. The earlier systems with analog circuits cannot be modified that easily.

So, we obtained the base pulse width from the table but as the operating conditions of automotive engines are



Idle speed is more complicated: the throttle is closed, so there has to be a bypass to let the engine receive fuel to run. In simpler systems this bypass is constant (but manually adjustable to set the correct idle speed, providing a fixed amount of air, although the computer can decide on a varying amount of fuel to be injected. Later systems generally use a controlling device changing the cross section of the bypass, regulating the amount of air coming through (these systems often have no facility to adjust the idle speed, the computer knows the correct revolution and maintains it without any help from mechanical devices). The controlling device


hardly ideal for any reasonable amount of time, several corrections have to be applied. Our air flow meter measures the volume of the air but we would need to know the mass of the air to calculate the required lambda ratio—remember, colder air is denser, thus the same volume contains more gas, requiring more fuel to provide the same mixture. To accomplish this, the injection system uses an **air temperature sensor (ATS)**—although on some systems it measures not the air but the fuel-air mixture—and lengthens the injector pulse width according to this input (except for the case of the airflow meter using a heated wire, this one takes the air temperature into account automatically, consequently, there is no need for correction).


It is not only the external circumstances that require special consideration. While most of the time an engine works under partial load, so it makes sense to spare fuel by basing on a relatively leaner mixture across this range of operation, cold start and warm-up, modest deceleration and fully depressed throttle, idle speed all require different treatment.

The position of the throttle pedal is communicated to the computer by a **throttle position switch (TS)** or **throttle potentiometer (TP)**. These devices signal both fully open and fully closed (idling) throttle positions. When the pedal is fully depressed, the computer makes the mixture richer to provide good acceleration performance.

can either be an **idle speed control valve (ISCV)** or an **idle control stepper motor (ICSM)**. The first one can only open or close the idle bypass, so any regulation must be done by rapidly opening and closing it by the computer, the second one can gradually change the bypass, hence fine tuning is easier and smoother.

Just like the choke on carburetors, there is a complete subsystem dealing with cold start and warm up, as the requirements under such circumstances are so different from the normal operation that they cannot be fulfilled by the regular control. The ECU monitors the **ignition key switch** to learn when the engine is started, then looks for the input from the **coolant temperature sensor (CTS)** to





lay depends heavily on the voltage the injectors are fed with. The same pulse width would result in shorter opening time, hence less fuel injected if the battery voltage drops below nominal (which is often the case on cold mornings). The injection computer therefore has to sense the battery voltage and to lengthen the injector pulse width if necessary.

The final, total pulse width (also called *injector duty cycle*) is calculated by summing up all these values received: the base pulse width from the RPM/AFS table lookup, the various correction factors based on the temperature sensors, throttle position and the like, plus finally, the voltage correction.

see whether this is a cold start or a warm one. If the coolant fluid is measured cold, a special warm-up sequence will be started.

The engine needs significantly more fuel, a richer mixture during this period. This extra fuel is used for two purposes: first, part of the fuel injected is condensed on the cold walls of the engine, second, to ensure better lubrication, the engine should run at an elevated revolution during this period.

There are two ways to provide more fuel: through the usual injectors, making the computer inject more gas than normal, or by using an additional **cold start injector (CSV)**—there is only one such injector even in multipoint systems. This injector is fed through a **temperature-timer switch**, protruding into the coolant just like the CTS, plus it is heated by its own electric heater. The injector operates as long as the ignition key is in the starting position but its behavior later on is governed by the timer switch. The colder the engine initially is, the longer it stays closed to let the cold start injector do its job. In a warm engine (above 40 °C) it does not close at all.

Without a cold start injector, the computer itself adds about 50% extra fuel initially and drops this surplus to about 25% until the end of a 30-second time period.

From that point, the surplus is dictated by the warming of the engine, communicated by the CTS to the computer. EFI systems without an idle speed control device often use an electromechanical **auxiliary air valve (AAV)**. This valve, which is fully open when the engine is still cold but will close gradually as it warms up, lets an additional amount of air measured by the AFS pass through the system. Because it is measured, it tricks the computer into providing more fuel. The valve is heated by its own heating element as well as the engine, thus it closes shortly.

The injectors are electrovalves. As with any electromagnet, there is a small time delay between the arrival of the control signal and the actual opening of the valve due to the build-up of electromagnetic fields. The length of this de-

As the computer has already calculated the exact amount of fuel to be injected, there is only one task left: actually injecting it. There are two possible ways: to inject the fuel into the common part of the inlet, still before the throttle butterfly, or to inject them close to the inlet valves, individually to each cylinder. Depending on the solution chosen, the system will be called **monopoint** or **multipoint**. Monopoint fuel injection requires a single common injector; the smaller cost and simpler setup makes it more common on smaller engines (in the case of Citroëns, the 1380 ccm ones). In all cases, the computer actually calculates the half of the fuel amount required as it will be injected in two installments, once for each revolution of the engine.

The injectors of the multipoint system can be operated simultaneously or individually. Previous Citroëns on the road today still use simultaneous operation. Individual cylinder injection, however, holds great potential—just to name one, some of the cylinders of a larger engine can be temporarily shut off by cutting off their fuel supply if the car is operating at partial load, saving a considerable amount of fuel—, so we are sure to meet this sort of fuel injection systems in the future.

All systems—regardless of the number of injectors—use a similar fuel supply layout. The fuel is drawn from the **tank** by a continuously operating **fuel pump**, transported via a filter to the injectors, then back to the tank. There is a **pressure regulator** in the circuit as well to keep the pressure of the fuel at a constant pressure above that in the inlet manifold (this regulator is a separate unit on multipoint systems while integrated into the injector on monopoint ones). As the pressure difference between the two sides of the injectors are constant, the amount of fuel injected depends solely on the opening time of the injectors. The pressure used in contemporary EFI systems is 3 to 5 bars.

This is practically all there is to it, there are only a couple of safety and economy features in addition. If the engine revolution exceeds a certain limit (between 1,200 and 1,500 usually) and the throttle is closed—this is called decel-

eration—, the momentum of the car is sufficient to rotate the engine through the wheels. To save fuel, the injection is cut off. As soon as the engine speed drops below the limit or the throttle is opened, the injection is reintroduced—supposedly smoothly and gradually, however, many drivers complain about some jerkiness.

To avoid prolonged operation at revolutions exceeding the specification of the engine, the injection is cut off above a maximum engine speed (6,000-7,000 rpm, depending on the engine). And finally, to avoid the hazard of fire in a crash and the fuel squirting from the injection system with the engine stopped or possibly destroyed, the relay of the injectors is controlled by the ECU, allowing fuel injection only when the ignition (or the signal of the corresponding sensor) is present.

Who will light our fire?

Models with simpler fuel injection have traditional (electronic) ignition systems which are practically equivalent to the solution used on cars with carburetors.

The **distributor** has two purposes: generating the driving signal for the ignition system and to distribute the high voltage to the four cylinders in turn. This two parts inside the distributor are electrically separate but mechanically coupled—both are driven by the camshaft to keep them in sync with the strokes of the engine.

The ignition signal thus starts from the distributor. A magnetic induction sensor (consisting of a rotating four-sided magnet and a pick-up coil) sends a pulse to the ignition module at each firing point. This pulse will be switched to the **ignition coil** (an autotransformer; *auto* here does not mean that it is manufactured for automotive use, autotransformers have their primary and secondary coils connected) by a power transistor inside the module. The current change in the primary coil induces very high voltage spikes in the secondary circuit. These spikes then go back to the HT part of the distributor which in turn sends them to the **spark plug** of the actual cylinder requiring the spark.

It takes some time for the spark to ignite the fuel-air mixture inside the combustion chamber: this means that the spark has to arrive slightly before the piston reaches its top position (top dead center, TDC), so that it will receive the downward force of the detonation in the right moment. However, as the engine speed increases, so does the speed of the piston or the distance it travels during a given period of time. Therefore, the exact time of the spark has to be advanced as the revolution increases. Traditional systems do this by adding a vacuum line connecting the inlet manifold to the distributor. As the

vacuum increases with the engine revolution, its sucking force rotates the inner part of the distributor slightly away from its original position, causing all its timing devices switch earlier, as required by the value of the timing advance.

Clever systems can get away without a distributor: some CXs have such an ignition setup. This systems has two ignition coils, both serving two spark plugs at the same time. These two spark plugs belong to cylinders whose pistons move in unison: one is compressing, the other exhausting. Although both plugs generate sparks at the same time, the one in the exhausting cylinder will be wasted.


Two birds with one stone

We made the ignition seem too simple in the previous section. While it works as described, there are many factors to be considered if we want to build a modern ignition system. For instance, the timing advance depends not only on engine speed but on many other factors as well: engine load, engine temperature and to some extent, the air temperature.

Just like the carburetor was not really good at deciding the amount of fuel required by the engine, the traditional ignition is similarly not perfect in estimating the timing advance and other characteristics of the sparks needed. An electronic system similar to the one used for fuel injection shows clear advantages over any earlier system.

And as they use about the same sensors and rely on each other, what could be more logical than to integrate them into a common system, elegantly called an **engine management system**?

If we compare the schematics of the corresponding EFI and EMS systems, they look almost the same. There are two notable differences: the small arrow on the line connecting the ECU to the distributor has changed its direction and a new sensor, a **crank angle sensor (CAS)** has appeared.



Both changes have to do with the fact that the enhanced system, whose new task is to generate the ignition signals as well, cannot at the same time build on them as inputs. This new sensor—practically a replacement for the induction magnet in the distributor of earlier systems—informs the computer of both engine speed and camshaft position.

The flywheel has steel pins set into its periphery. As it rotates, the inductive magnet of the CAS sends pulses to the computer. Two of the pins are missing and this hole passes before the sensor just as the first piston reaches its TDC position. The missing pins cause a variance in the sensor output that can be read by the ECU easily.

The rest is the same: the base pulse width is calculated based on the CAS and AFS/MAP sensors. The correction factors—air temperature, idle or full load, starting, warming up, battery voltage—sum up into an additional pulse width. Besides, the same input signals (AFS, CAS, CTS and TS/TP) are used for another lookup in a table, yielding the correct dwell time and timing advance for the ignition. The dwell period remains practically constant but the duty cycle varies with the changing engine speed. The ignition signal is amplified and sent to a distributor containing only secondary HT components: it does not create the ignition signal only routes the HT current to each spark plug in firing order.


Some systems also have a **knock sensor (KS)**, sensing the engine vibration associated with pre-ignition (so-called pinking). If this occurs, the ignition timing is retarded to avoid engine damage.

Think green

As we saw, fuel injection and engine management systems are capable of determining the ideal amount of fuel to be injected, depending on the conditions of operation and several other factors in the engine. It is capable of deciding on lean mixture for general, partial load to save fuel, or on rich mixture when performance considerations call for this.




Unfortunately, this is not what such systems are used for today. With the proliferation of catalytic converters, the only concern of our systems is the welfare of the converter.

Ideal combustion would not generate polluting materials in the exhaust gas. Fuel is a mixture of various hydrocarbons (C_nH_m), which when burned together with the oxygen (O_2) of the air, should transform to carbon-dioxide (CO_2) and water vapor (H_2O). However, combustion is never ideal, besides, fuel contains many additives: the exhaust gas, in addition to the products mentioned, has various byproducts as well, some of them toxic: carbon-monoxide (CO), various unburned hydrocarbons (C_nH_m), nitrogen-oxides (NO_x) and lead (Pb) in various substances coming from the anti-knock additives found in the fuel.




The relative amount of these byproducts depend on the lambda ratio of the air-fuel mixture burned. As shown on the diagram, a value between 1.2 and 1.3 would give a relatively low percentage of toxic byproducts while, as we can recall, being a lean mixture would be in the right direction towards fuel economy.

By using platinum (Pt) or rhodium (Rh) as a catalyst—a catalyst is a substance whose presence is required to enable (or to boost) a chemical transformation while it does not take part in the process itself, remaining intact—the following processes can be carried out:



These precious metals are applied in a very thin layer to the surface of a porous ceramic body with thousands of holes to make the surface contacting the exhaust gases much greater. Actually, a converter does not contain more than 2 or 3 grams of these metals.

If you compare this diagram with the previous one, you will see that the real gain is the suppression of nitrogen-oxides. CO and C_nH_m will be reduced as well, although to a much lesser extent. Nevertheless, the overall reduction in polluting byproducts is quite high, amounting up to 90 percent. Lead substances are not considered as lead must not reach the converter anyway, it would clog the fine pores of the converter in no time. The fuel used in cars equipped with a catalytic converter has to be completely free of lead.



But there is something of even greater consequence depicted on the diagram: to keep the amount of pollutants down, the lambda has to be kept inside a very small value range, practically at $\lambda=1$ all the time. If the lambda drops just a fraction below 1, the CO emission rises sharply, while a small step above 1 skyrockets the NO_x emission. The main task of the fuel injection is therefore to ensure that the air-fuel mixture sticks to the stoichiometric ratio all the time. This means higher consumption than the one of a car with fuel injection without a converter to start with.

There are situations where this lambda cannot be observed. A cold engine will simply stall without a much richer mixture, thus the cold start mechanism does not obey the lambda control. The catalytic converter does not work at all below 250 °C, so this is not a significant compromise (its normal operating temperature is 400 to 800 °C, above 800 °C is already harmful; unburned fuel getting into the exhaust and detonating inside the converter could cause overheating, thus ignition and similar problems has to be rectified as soon as possible in catalytic cars).

Dynamic acceleration (full throttle) is also something not observing the welfare of the converter. Reducing pollution might be a noble cause but to be able to end an overtaking is even more important...

The system uses an **oxygen sensor (OS)**, also called lambda sensor) which measures the oxygen content of the exhaust gas. It is located between the engine exhaust and the catalytic converter. Similarly to the converter, it is not functional below 300 °C, hence it has its own heating element to make it reach its operating temperature faster.

The computer uses the input from this sensor to keep the mixture injected always as close to $\lambda=1$ as possible. If the sensor is still too cold to give accurate input, the computer can ignore it safely.

Diesel engines

Diesel oil has been a contender to gasoline for many decades. Earlier diesel engines were not refined enough to win the hearts of many drivers but recent advances in technology made these engines not only worthy competitors in all areas but in some features—fuel economy or low end torque, to name just two—even exceeding the characteristics of their gasoline counterparts. And in addition to the general technological advances, Citroën's diesel engines have a widely accepted reputation—even among people blaming the quirkiness of its suspension or other features—of being excellent and robust.


As it is widely known, diesel engines have no ignition to initiate their internal combustion, they rely on the self-combustion of the diesel oil entering into a cylinder filled with hot air. Due to this principle of operation, the supply of the fuel has to comply with much more demanding requirements than it is necessary in the case of gasoline engines.

Unlike in the gasoline engine, not a mixture but air enters into the cylinders via the inlet valves. During the adiabatic compression all the energy absorbed is used to increase the temperature of the gas. The small droplets of fuel will be injected at high velocity near the end of the compression stroke into this heated gas still in motion. As they start to evaporate, they form a combustible mixture with the air present which self-ignites at around 800 °C.

This self-ignition, however, is not instantaneous. The longer the delay between the start of the injection and the actual ignition (which depends on the chemical quality of the

diesel oil, indicated by the *cetane number*), the more fuel will enter the cylinder, leading to harsher combustion, with the characteristic knocking sound. Only with the careful harmonization of all aspects—beginning of injection, the distribution of the amount injected in time, the mixing of the fuel and air—can the combustion be kept at optimal level.

Small diesel engines suitable for cars were made possible by a modification to the basic principle, that allowed these stringent parameters to be considerably relaxed. It includes a separate swirl chamber connected to the cylinder via a restrictor orifice. The air compressed by the piston in the cylinder enters this chamber through the orifice, starting to swirl intensively. The fuel will then be injected into this swirl, and the starting ignition propels the fuel-air mixture still incompletely burned into the cylinder where it will mix with the air, continue and finish the combustion process. Using a prechamber results in smaller ignition delay, softer combustion, with less noise and physical strain on the engine parts, but introduces some loss of energy because of the current of air having to pass between the chambers. Citroën engines of this type use a tangentially connected spherical prechamber.



As diesel engine evolution continued, better simulation and modeling techniques became available, which, together with the improvements in fuel injection technology, lessened or removed the problems initially solved by the introduction of the prechamber. The direct injection engines of today have no prechamber, instead, the piston has a specially formed swirl area embedded in its face.

Mechanical injection

Although the basic principles of fuel injection are similar to what we have already discussed for gasoline engines, there are some notable differences. First of all, diesel engines operate without restricting the amount of air entering the engine: there is no throttle, the only means of regulating the engine is to vary the amount of fuel injected.

The fuel is injected into the engine, creating a combustible mixture in the same place it is going to be burned. Because the forming of this mixture results in its self-combustion, the diesel injection system is, in essence, an ignition control system. Unlike on the gasoline engine, fuel injection and ignition cannot be separated in a diesel engine.

The complete mechanical injection system is built into a single unit which can be divided into five individual—although interconnected—subsystems:


- ✦ a **low pressure fuel pump** to deliver the fuel for the rest of the injection system;
- ✦ a **high pressure pump and distributor** that routes the fuel to the appropriate cylinders in firing order (similar in purpose to the distributor on gasoline en-

gines) and generates the high pressure needed for the injection as well;

- ✦ a **regulator** that determines the amount of fuel to be injected in relation to the engine speed, modified by additional factors like idle speed, cold starting, full load, etc.;
- ✦ an **injection adjuster** to compensate for the higher engine speed by advancing the start time of the injection;
- ✦ a **fuel stop valve** to cut off the fuel supply when the ignition has been switched off.

The diesel fuel is drawn—through a filter—from the tank by the **low pressure pump ①** operated by the engine. A pressure regulating valve **②** ensures that the fuel pressure will not exceed a preset limit; when the pressure reaches this value, the valve opens and lets the fuel flow back to the primary side of the pump.

The **piston ⑥** of the high pressure part is driven through a coupling **④** consisting of a cam disc and four cam rollers. The piston rotates together with the shaft coming from the



engine but the coupling adds a horizontal, alternating movement as well: for each turn, the shaft and the piston 6 performs four push-pull cycles.

It is the pushing movement of this piston 6 that creates the high pressure and sends the fuel to the injectors. The fuel, provided by the pump 1 arrives through the fuel stop electro-valve 17, which is constantly open while the ignition switch is on but cuts the fuel path when it is turned off.

First, the piston 6 is pulled back by the coupling 4, letting the fuel enter the chamber and the longitudinal bore inside the piston. As the side outlets are blocked by the regulator collar 5, the fuel stays inside the chamber (*phase 1*).

In the next phase, the piston rotates and closes the ingress of fuel from the stop valve 17. On the other side of the piston, the high pressure outlet opens but as the fuel is not yet under pressure, it will stay in the chamber.

In *phase 3* the piston is energetically pushed by the cam disc and rollers of the coupling 4, injecting the fuel stored in the chamber into the output line with a significant force.

As the piston 6 moves to the right, at some point the side outlets will emerge from under the regulator collar 5—the fuel injection into the real output will stop immediately, and the rest of the fuel stored in the chamber will leave through this path of lesser resistance. This is *phase 4*, the end of the injection cycle.

Actually, this operation is repeated four times for each revolution of the incoming shaft. There are four high pressure outlets radially around the piston, each serving a given cylinder. As the outlet slot 19 of the piston turns around, it allows only one of the outlets to receive the fuel.


The pressure valves 7 serve to drop the pressure in the injector lines once the injection cycle is over. To reduce the cavitation caused by the pressure waves generated by the rapid closing of the injector valves, a ball valve minimizing the back flow is also used.


The length of phase 3, thus the amount of fuel injected depends on the position of the collar 5. If it is pushed to the right, it will cover the side outlets for a longer time, resulting in a longer injection phase, and vice versa. If it stays in the leftmost position, no fuel will be injected at all.

And this is exactly what the **regulator part** does: it moves this collar 5 to the left and to the right, as the actual requirements dictate. The lever 9 attached to the collar is rotated around its pivot by several contributing forces. The two main inputs are the position of the **accelerator pedal** as communicated through a regulator spring 12 and the actual **engine speed**, driving a centrifugal device 8 via a pair of gears 3. The higher the engine speed, the more the shaft 20 protrudes to the right, pushing on the lever 10.

When the engine is being started, the centrifugal device 8 and the shaft 20 are in their neutral position. The starting lever 10—pushed into its starting position by a spring 11—sets the position of the collar 5 to supply the amount of fuel needed for the starting.

As the engine starts to rotate, a relatively low speed will already generate a large enough force in the centrifugal device 8 to push the shaft 20 and overcome the force of the





rather weak spring 11. This will rotate the lever 10, moving the collar 5 to the left, setting the amount of fuel required for idling. The accelerator pedal is in the idle position as well, dictated by the adjustment screw 14. The idle spring 13 keeps the regulator in equilibrium.

Normally, the amount of fuel will be regulated by the position of the pedal as both springs 11 and 13 are fully compressed and do not take an active part in the process. When the driver pushes on the pedal, the regulating spring 12 stretches, both levers 9 and 10 rotate and move the collar 5 to the right, to allow the maximum amount of fuel to be injected. As the actual engine speed catches up, the centrifugal device 8 opens up, pushing the shaft 20 to the right, countering the previous force, gradually returning the collar 5 towards the no fuel position, until the point is reached where the amount of fuel injected maintains the equilibrium. When the driver releases the pedal, the inverse of this process takes place. During deceleration—pedal at idle, engine rotated by the momentum of the car—the fuel is cut off completely.

Without such regulation, if enough fuel is provided to overcome the engine load, it would continue accelerating until self-destruction (this is called *engine runaway*). Speed regulation is a feedback mechanism comparing the actual speed of the engine to the one dictated by the gas pedal and modifies the amount of fuel as necessary. If either the engine speed changes (because of varying load, going over a hill, for instance) or the driver modifies the position of the accelerator pedal, the regulation kicks in, adding more or less fuel, until a new equilibrium is reached. If the engine is powerful enough to cope with the load, keeping the pedal in a constant position means constant cruising speed in a diesel car; gasoline vehicles need speed regulated fly-by-wire systems or cruise controls to achieve the same.

The excess fuel will finally leave the pump unit through an overflow valve 18, flowing back to the fuel tank.

Something needs to be corrected...

The chemistry involved in the combustion dictates some parameters of fuel injection, the most important being the **smoke limit**, the maximum amount of fuel injected into a given amount of air, that results in combustion without resulting in soot particles. Although gasoline engines also have this limit, they normally operate with a constant fuel to air mixture that automatically places the amount of fuel below this critical limit. Diesel engines, in contrast, operate with a variable fuel to air mixture, using this very variation for power regulation. With diesel fuel observing the smoke limit is a much stricter task because once soot starts to develop, this changes the character of the combustion itself, resulting in a sudden and huge increase in the amount of particulates—a bit like a chain reaction.


Because the maximum amount of fuel injected depends on how far the lever 10 is allowed to rotate counter-clockwise, the inability of the pump to inject too much fuel, thereby crossing the smoke limit, is insured by an end stop 2 for this lever. This very basic means of smoke limit correction, adjusted for worst case conditions, was developed further on turbocharged engines, and still further on electronically controlled injection systems.

Timing is of enormous importance in a diesel engine. During the stroke of combustion, several events take place in close succession: the fuel injection system starts its delivery, then the fuel is actually injected (the time elapsed between these two is the *injection delay*), slightly later the fuel will self-ignite (this delay is the *ignition delay*), then the injection will stop but the combustion is still raging, first reaching its maximum, then dying away slowly (on the scale of milliseconds, that is).

Just like in a gasoline engine, the ignition delay remains constant while the engine speed changes. The fuel has to ignite before the piston passes its TDC position, but with the increasing engine speed, the distance the piston travels during a given period of time becomes longer. Therefore, the injection has to be advanced in time to catch the piston still in time. The **injection adjuster** 15 feeds on the fuel pressure provided by the pump 1, proportional to the engine speed.


This will move the piston, which in turn, through the levers, modifies the relative position of the cam rollers to the cam disc inside the coupling 4, increasing or decreasing the phase difference between the revolutions of the engine and the rotating-alternating movement of the distributor piston 6.

Some engines also have additional minor **correction mechanisms** 16 that modify the idle speed and timing depending on engine temperature, to provide better cold start performance. The engine temperature is measured indirectly, through the coolant acting on cylinder and piston-like elements filled with paraffin. As the paraffin expands or contracts as the coolant temperature dictates, the transformed mechanical movement, coupled through cables to two movable end stops for both the lever 9 and the injection adjuster 15, modifies the idle speed and the injection timing of the engine. Because correct timing depends on temperature, the corrections, although relatively slight, in-



sure that the amount of fuel injected as well as the timing provide better combustion and lower pollution when the engine is started and operated at low temperatures. They do not have any effect once the engine reaches the normal operating temperature.

Now that the correct amount of fuel is carefully determined and the necessary high pressure generated by the pump, it has to be injected into the swirl chamber. The pressurized fuel entering the **injector** through a filter ① tries to press the piston ② upwards but a spring ③ counters this force. As soon as the pressure exceeds the force of the spring (which can be adjusted by placing appropriately sized shims behind it), the piston jumps up and the fuel rushes into the swirl chamber




through the small orifice now opened. After the injection pump closes its pressure valve at the end of the injection period, the spring ③ pushes the piston ② back, closing the orifice until the next injection cycle.

Each swirl chamber has its own **glow plug** whose only purpose is to heat up the chamber in cold weather. They start to glow when the ignition key is turned into the first position and stay glowing for some time afterwards unless the starting was unsuccessful.

Turbo

More power requires more fuel. An efficient way to boost the performance is to provide both more air and fuel to the engine. The exhaust gases rushing out from the engine waste a great deal of energy; a **turbocharger** ④ spun



by the exhaust flow taps into this source of energy to provide added pressure in the air inlet. Diesel engines are particularly well suited for turbocharging. Gasoline engines may not have the inlet pressure raised too much because the air and fuel mixture may subsequently self-ignite when it is not supposed to, and instead of burning controllably, detonate. In a diesel such a situation is not possible because the fuel is injected only when combustion should actually happen in the first place. As a result, relatively high inlet pressures can be used, considerably improving the power output of a diesel engine, and with proper attention to the subtleties of the design, engine efficiency and fuel consumption.

On its own, once the amount and pressure in the exhaust manifold reaches a level high enough to power it, with the engine fully loaded, the turbine would spin proportionally to engine speed squared, because both the pressure and the volume of the air pumped into the engine are increasing.

Because the engine is required to deliver as much torque as possible at the widest possible range of engine revolution, the requirements on the turbine are somewhat contradictory. If the turbo is made very small and light, it will spin up very quickly due to its low mass and inertia, ensuring its full benefit already at low rpms. However, with a moderate


increase in engine speed, the rotational speed of the turbine (note the quadratic relationship) would become excessively high. When the turbine blade speed approaches the speed of sound, a supersonic wave effect occurs that can abruptly leave it without any load, at which point runaway would occur, resulting in severe damage to the turbine.

On the other hand, if the turbine was dimensioned so that even at the highest engine speed it is still operating within safe limits, it would not be useful at all in the middle range where the engine is most often used. A compromise can be achieved using an overpressure valve, the **wastegate valve** ⑤. The turbo pressure is constantly monitored by this valve opening above a set pressure limit, letting the exhaust escape through a bypass. This avoids turbo runaway by making the turbo rotational speed proportional to that of the engine, once the limit pressure is reached. This way the quick spin-up resulting from the quadratic relationship can be preserved while the turbocharging effect is extended over a significant percentage of the usable engine speed range—typically the higher 70-80%. But it comes at a price: because of the simplicity of such a regulation, the limit pressure is dictated by the maximum turbine speed, which is usually calculated for maximum engine speed plus a safety margin. The maximum pressure is already reached at lower engine and turbine speeds, where the turbine could conceivably still provide more pressure because of a lesser demand for air volume. Although with a simple wastegate a certain amount of the turbocharging potential is lost, the increase in power output is still substantial.

Citroën is a pioneer in implementing variable wastegate limit pressure using a controllable wastegate valve, to tap into this previously unused turbo potential.

Essentially, a turbocharged diesel engine runs in two different modes: atmospheric pressure or turbo-charged. The atmospheric pressure mode prevails while the exhaust gas produced is not yet sufficient to power the turbine (below a given engine speed and load). Once this limit is crossed and the turbine starts generating higher than atmospheric pressure, the engine is running in turbocharged mode.

The injection pump regulator needs to know about the changes in the inlet pressure, because those changes mean differences in the amount of air entering the engine. And this also means that the upper limit of fuel injected needs to be changed correspondingly. These injection systems are tuned for the turbo producing the rated waste pressure (also known as full boost).



However, the amount of fuel injected during the atmospheric mode of the engine—before the turbo kicks in—has to be reduced in order to avoid crossing the smoke limit. The turbo pressure drives a limiter in the injection pump: with the increasing pressure the piston ②① moves down. Its varying diameter forces the lever ②② rotate around its pivot, which then acts as a stop to limit the allowed range of operation of the regulator lever ⑨, limiting the amount of fuel to be injected.

Intercooler

Towards a cleaner world

Exhaust Gas Recycling (EGR) systems were used—depending on the market—as add-on units. An electronic unit measuring the coolant temperature and the position of the gas pedal control on the pump (with a potentiometer fitted to the top of the control lever) controls a valve which lets part of the exhaust gas get back into the inlet.

Post-glowing is also used as a pollution reducing mechanism. A definite post-glow phase, lasting for up to minutes is usually controlled by a combination of a timer and the engine coolant temperature: either the timeout of 4 minutes runs out or the engine reaches 50 °C. An additional mechanism prevents post-glowing if the engine was not actually started.

Electronic Diesel Control

Just like it is the case with gasoline engines and carburetors, a mechanical device—even one as complicated as a diesel injection pump—cannot match the versatility and sensibility of a microcomputer coupled with various sensors, applying sophisticated rules to regulate the whole process of fuel injection.

The only input a mechanical pump can measure is the engine speed. The amount of air entering into the engine, unfortunately, is far from being proportional to engine speed, and the turbo or the intercooler disturbs this relationship even further. As the injection always has to inject less fuel than the amount which would already generate smoke, the mechanical pump—capable only of a crude approximation of what is actually going on in the engine—wastes a significant amount of air, just to be of the safe side.

The satisfactory combustion in diesel engines relies on the exhaust as well—if this is plugged up, more of the exhaust gases stay in the cylinder, allowing less fresh air to enter. A mechanically controlled injection pump has no feedback from the engine (except for the engine speed)—it will simply pump too much fuel into the engine, resulting in black smoke. An electronically controlled injection pump, on the other hand, can tell how much air has actually entered by using a sensor (although only the latest systems use such a sensor).

There are also other factors never considered by a mechanical system. The details of the combustion process depend heavily on the chemical characteristics of the fuel. The ignition delay, as we have already seen, depends on the cetane number of the diesel oil. In spite of the fact that correct timing has a paramount influence on the performance and the low pollutant level of a diesel engine, the mechanical system can have no information about this very important input factor. Less essential but still important is the temperature of the incoming air. With measuring all the circum-


stances and conditions in and around the engine (air, engine and fuel temperatures), the injection system can achieve better characteristics, lower fuel consumption and less pollution.

All in all, the electronically controlled injection pump not only adds precision to the injection process as its gasoline counterpart does but introduces completely new methods of regulation; therefore it represents a much larger leap forwards than fuel injection in gasoline engines. In spite of this, it is quite similar to its mechanical predecessor. From the five subparts, four remain practically the same, only the **regulator** is replaced with a simple electromagnetic actuator that changes the position of the same regulator collar ❸ as in the mechanical pump, in order to regulate the amount of fuel to be injected.

The real advantage over the former, mechanical pumps is that an electronic device, a small microcomputer can handle any complex relationship between the input values and the required output. With mechanical systems, only simple correction rules are possible, and as the rules get more complicated, the mechanics quickly becomes unfeasible. In contrast to this, the ECU just have to store a set of characteristic curves digitized into lookup tables, describing the amount of fuel to be injected using three parameters: **engine speed** (measured by a flywheel inductive magnet), **coolant temperature** (measured by a sensor protruding into the coolant liquid), **air temperature** (measured by a sensor in the air inlet).

The newer HDi engines use an **air mass sensor** using a heated platinum wire (as that mentioned on page 6). Having the exact amount of air to enter the engine, these latest EDC systems can deliver true closed loop regulation.

A potentiometer attached to the **accelerator pedal** sends information about the pedal position to the computer. This signal is used as the main input, conveying the in-



tentions of the driver. The ECU uses this sensor to learn about special conditions like idle speed or full load as well.

Air temperature is measured by a sensor in the inlet manifold (but if the air mass is measured by a heated platinum wire sensor, this already provides the necessary air temperature correction, thus there is no need for an additional sensor).

The ECU stores the basic engine characteristics, the intrinsic relationship between the air intake and the engine speed (plus the manifold pressure if a turbo is fitted). The values obtained from this table are corrected according to the inputs of the various sensors, in order to arrive at a basic timing and smoke limit value. The actual amount of fuel injected and the accurate timing are a function of these results and the position of the accelerator pedal.

The final amount of fuel calculated will be used to control the electric **actuator** ⑨ which—by moving a lever ⑩—changes the position of the regulating collar ⑤. To ensure the necessary precision, the factual position is reported back to the computer using a potentiometer.

As we have already mentioned, the exact timing of the injection is of utmost importance in a diesel engine. The electronic system uses a **needle movement sensor** built into one of the injectors (the other are assumed to work completely simultaneously) notifying the computer about the precise time of the beginning of the injection. Should there be any time difference between the factual and designated opening times, the electro-valve ⑭ of the **injection adjuster** ⑮ will receive a correction signal until the difference disappears. If the electro-valve is completely open, the injection start will be delayed, if it is closed, the start time will be advanced. To achieve the timing required, the valve is driven with a modulated pulse signal, with the duty cycle (on-off ratio) determined by the ECU.

The input from this sensor is also used for compensating calculations on the amount of fuel injected, and to provide the on-board computer with the exact amount of fuel used up so that it can calculate the momentary and average consumption.

The computer has extensive self-diagnostic functionality. Many sensors can be substituted with standard input values in case of a failure (serious errors will light up the diagnostic warning light on the dashboard). Some sensors can even be simulated using other sensors—for instance, the role of a failing engine speed sensor might be filled in by the signal generated from the needle movement sensor.

As there is no standalone ignition in a diesel engine, the only way to stop it is to cut off the fuel supply. The mechanical default position of the actuator ⑨ is the position where no fuel enters the injectors at all; this is where it returns when the computer receives no more voltage from the battery, the ignition switch having turned off.

As it has already been mentioned, the inlet pressure is one of the principal EDC parameters for a turbocharged engine. Later Citroën turbocharged diesels—starting with the 2.5 TD engine of the XM—pioneered variable turbo pressure technology. The wastegate on these turbines has several actuators, fed with the turbo pressure through electric valves. The ECU, based on the relevant engine operation parameters obtained from the sensors, controls these actuators in various combinations, providing a selection of two or three different wastegate limit pressures. This lets the sys-

tem ease the compromise between the turbo pressure and turbine speed: the pressure is kept at the usual value for higher engine speeds (limited by the maximum turbine speed) but is allowed to go higher than that in the middle rpm ranges, adding a significant amount of torque in the range where it is most needed.

Green versus Black

Diesel oil, just like gasoline, is a mixture of various hydrocarbons (C_nH_m), and burned together with the oxygen (O_2) of the air, transforms to carbon-dioxide (CO_2) and water vapor (H_2O). However, as the combustion is never ideal, the exhaust gas also contains various byproduct gases: carbon-monoxide (CO), various unburned hydrocarbons (C_nH_m), nitrogen-oxides (NO_x). The relatively high lambda value a diesel engine is operating with reduces the hydrocarbon and carbon-monoxide content to 10–15%, and the amount of nitrogen-oxides to 30–35% of the corresponding figures measured in gasoline engines without a catalytic converter. The sulphur content of the fuel—drastically reduced during the recent decades—is responsible for the emission of sulphur-dioxide (SO_2) and sulphuric acid (H_2SO_4).

Conversely, these engines emit 10–20 times more particulates—or black soot—than gasoline engines. These are unburned or incompletely burned hydrocarbons attached to large particles of carbon. These substances are mainly aldehydes and aromatic hydrocarbons; while the first only smells bad, the second is highly carcinogenic.

The much higher amount of particulates is due to the different combustion process. The various aspects of mixture formation, ignition and burning occur simultaneously, they are not independent but influence each other. The distribution of fuel is not homogenous inside the cylinder, in zones where the fuel is richer the combustion only takes place near the outer perimeter of the tiny fuel droplets, producing elemental carbon. If this carbon will not be burned later because of insufficient mixing, local oxygen shortage (large fuel droplets due to insufficient fuel atomization, caused by worn injectors) or the combustion stopping in cooler zones inside the cylinder, it will appear as soot in the exhaust. The diameter of these small particles is between 0.01 and 10 μm , the majority being under 1 μm . Keeping the amount of fuel injected below the *smoke limit*—the lambda value where the particulate generation starts to rise extremely—is essential.

Similarly to gasoline engines, the exhaust gas can be post-processed to reduce the amount of pollutants even further. There are two different devices that can be used:

❑ **Soot burning filter:** as the diesel engine always operates with excess air (its lambda is above 1), there is enough oxygen in the exhaust gas to simply burn the carbon soot present. The burning filter is manufactured from ceramic materials that can withstand the resulting high temperatures (up to 1200 °C). As the diesel engine is very sensitive to excessive back pressure, the filter has to be able to self-regenerate. This is solved by the addition of organic metal substances.

❑ **Catalytic converter,** identical to the simpler ones used on gasoline engines before the proliferation of three-way, controlled converters. It reduces the carbon-monoxide and hydrocarbon content of the exhaust gas.

Diesel Direct Injection

I think that at this point, soot burning filters will have to be cut out of the PDF and put in at a similar ecological section under DI/HDI—since that is the only system that actually makes soot burning practical, and the only system that implements it.

Soot burning was experimented with a lot but was never made practical before HDI due to a too low exhaust temperature. The particle filter would need heating to a very high temperature and that was deemed to be too dangerous. Even with cerine additives, essentially, there would have to be a separate small burner to heat up the filter, which is again another system that can go wrong. HDI essentially integrates a burner by allowing post-injection, something that is simply impossible for injection systems derived from a classical pump due to the timing required. I think that for soot management it is enough to write that the smoke limit control is *vastly* improved by the better regulation of the EDC.

Other things like controlled swirl and multi-valve technology, also pioneered by Citroën (XM 2.1 TD!) should be mentioned. The catalytic converter section remains unchanged.

And, of course, there should be an "In addition to the pollution management implemented on mechanical injection systems" sentence somewhere in there, since proper cold start corrections and EGR are implemented in EDC units by default.

Suspension



A Suspension Primer

From the early days of the automobile—and even before, in the time of horse-drawn carts—it was already well known that the body of the car, housing both the passengers and the load, has to be decoupled from the unevenness of the road surface.

This isolation is much more than a question of comfort. The vertical force of the jolts caused by the repeating bumps and holes of the road surface are proportional to the square of the vehicle speed. With the high speeds we drive at today, this would result in unbearable shock to both people and the mechanical parts of the car. Jolts in the body also make it more difficult to control the vehicle.

Consequently, there has to be an elastic medium between the body and the wheels, however, the elasticity and other features of this suspension medium are governed by many, mostly contradicting factors.

The softer, more elastic the spring, the less the suspended body will be shaken by various jolts. For the sake of comfort, we would thus need the softest spring possible. Unfortunately, too soft a spring will collapse under a given weight, losing all its elasticity. The elasticity of the spring would need to be determined as a function of the weight carried but the weight is never constant: there is a wide range of possible load requirements for any car. On one hand, a hard suspension will not be sensitive to load variations but being hard, will not fulfill its designated purpose, either. A soft suspension, on the other hand, is comfortable but its behavior will change significantly on any load variation. To cope with this contradicting requirements, an elastic medium of decreasing flexibility would be required: such a spring will become harder as the weight to be carried increases.

When the spring is compressed under the weight of the load, it's not only its flexibility that changes. The spring deflects, causing the clearance between the car and the road surface decrease, although a constant clearance would be a prerequisite of stable handling and roadholding. At first sight, this pushes us towards harder springs: soft springs would result in excessive variations of vertical position — unless, of course, we can use some other mechanism to ensure a constant ground clearance.

In addition to the static change caused by load variations, the deflection of the spring is changing constantly and dynamically when the wheels roll on the road surface. The body of the vehicle dives, squats, rolls to left and right as the car goes over slopes, holes and bumps in the road, corners, accelerates or decelerates.

When a deflected spring is released again, the energy stored in it will be released but as there is no actual load for this energy, the elastic element, the mass of the suspension and the vehicle form an oscillatory system, causing a series of oscillations to occur instead of the spring simply returning to its neutral position.

Any vertical jolt would thus cause such oscillations: the upward ones are transmitted to the car body while the downward ones make the wheels bounce, losing contact with and adhesion to the road surface. The first is only discomforting, but the second is plainly dangerous. In addition, it's not only the spring that oscillates; the tires contain air which is a highly elastic spring medium. Oscillation in itself causes unwanted motion but when the corrugation of the road surface happens to coincide with the period of the suspension oscillations, it might lead to synchronous resonance, a detrimental situation leading to serious damages in the suspension elements.

Mass in motion can also be viewed as a source for kinetic energy; because of this, moving parts of the suspension are often reduced in weight to decrease this portion of the stored energy, and this in turn eases the requirements on the dampers as they have to dissipate less unwanted energy as heat. This solution, however, often shifts the frequency of the self-oscillation of the suspension upwards. Unfortunately, occupants are more sensitive to higher frequencies reducing comfort (mostly adding noise), so this is an area where compromise is needed.

Conventional suspension systems use a second element, a shock absorber to dampen these oscillations. The absorber uses friction to drain some of the energy stored in the spring in order to decrease the oscillations. Being an additional element presents new challenges: the characteristics of both the spring and the absorber have to be matched carefully to obtain any acceptable results. The absorber ought to be both soft and hard at the same time: a soft absorber suppresses the bumps of the road but does not decrease the oscillations satisfactorily while a hard absorber reduces the oscillations but lets the passengers feel the unevenness of the road too much. Due to this contradiction, conventionally suspended cars have no alternative but to find a compromise between the two, according to the intended purpose of the car: sport versions are harder but offer better roadholding, luxurious models sacrifice roadholding for increased comfort. This contradiction clearly calls for a unified component serving both as a spring and an absorber, harmonizing the requirements.

Hydropneumatic Suspension

As we saw, the ideal suspension would require elasticity decreasing with the load, constant ground clearance, shock absorbers integrated into the suspension—all these beyond the obvious independent suspension for all wheels. And this is exactly what Citroën's unique hydropneumatic suspension offers.

According to the Boyle–Mariotte formula defined in the 17th century, the pressure and the volume of a mass of gas are inversely proportional at a constant temperature. Therefore, by keeping the mass of the gas constant and changing the volume of its container, its pressure can be controlled (the usual pneumatic suspensions operate on the opposite principle: air is admitted or withdrawn from the system by compressors and exhaust valves, modifying its mass while keeping the volume constant).

The volume changes are controlled by hydraulics, a technology in widespread use in every branch of the industry. As liquids are non-compressible, any amount of liquid introduced at one end of a hydraulic line will appear immediately at the other end (this phenomenon was first formulated by Blaise Pascal). Using this principle, motion can be transmitted, multiplied or divided (according to the relative sizes of the operation cylinders), with velocity increased or decreased (using varying cross sections in the tubing), to any distance desired, over lines routed freely.

Hydraulics are immensely useful, very efficient, reliable, simple to use, and—due to their widespread deployment—relatively cheap. It is no wonder that it is used for many purposes even in the most conventional vehicles: shock absorbers, brake circuit and power assisted steering being the most trivial examples; however, Citroën is the only one to use it for the suspension.

The First Embodiment

DS



The Citroën DS, introduced at the 1955 Paris Motor Show, was radically different from any of its competitors on the market at that time: suspension, running gear, steering, brakes, clutch, body, aerodynamics were all unique, not only in details but in the main operating principles as well.

ID

The hydropneumatic spring-absorber unit uses an inert gas, nitrogen (colored blue on the illustrations) as its spring medium, resulting in very soft springing. The flexibility of the gas decreases as the increasing load compresses the suspension pistons, reducing the volume of the gas and adding to its pressure. The damping effect is obtained by forcing the fluid (colored in green) pass through a two-way restrictor unit between the cylinder and the sphere. This effect provides a very sensitive, fast and progressive damping to reduce any unwanted oscillations.

There are many great advantages to this hydropneumatic suspension. First, by adding or removing fluid from the suspension units (practically, by adjusting the length of the hydraulic strut), **ground clearance** can be kept **constant under any load variations**. Although this might not seem very important at first sight, it means that the suspension geometry is also constant—in other words, the handling of the car does not depend on the load.

The compressed gas has a variable spring effect, becoming harder as the load increases. This compensation for the increasing load keeps the **resonance frequency** of the suspension nearly constant. As a consequence, the same excitation in the suspension moves the same amount of fluid through the dampers regardless of load (which is not the case with conventional springs). The working range of the



dampers becomes much smaller and this fact makes the use of a simple damper element very effective.

This basically constant suspension resonance frequency also contributes to the consistent behaviour independent of the load. In essence, it ensures that both the road contact and the feeling transmitted to the driver remains always the same. This is something absolutely unique: all conventional suspensions have an optimum point around average load; when carrying more or fewer passengers or load than this average value, the handling characteristics change, not seldom so radically that the car becomes utterly dangerous to drive.

Another advantage is the limited but very useful **anti-dive behavior**: this is essential for efficient braking with a basically very soft suspension. The center of mass of the car moves much less than usual, hence the braking force is distributed more evenly. Manufacturers of cars with conventional suspension and braking only start to add brake force distributors to their vehicles these days. The first DS did have a force distributor but Citroën later realized that the suspension, with the addition of a single pipe, can fulfill its role entirely.

The height correction and the constant connection between the left and right side of the suspension has another important implication: lower difference in forces on the wheels. Coupled with variable damping this keeps the wheels **in contact with the road at all times**, which in turn maximizes the tractive forces on the tires—braking while turning still leaves the vehicle with the grip of all four wheels: this is essential for security in low adherence conditions, such as ice, snow, rain, mud.

GS

GSA

The steady connection between the sides requires an external management of body roll. Ideally, for any vertical movement of the car body, the two sides of the suspension should be connected, while for any movement that results in different displacements of each wheel, they should ideally be separate. This second movement can be viewed as a rotation around the longitudinal or transversal axis.

For instance, if the front wheels run into a pothole and the rear wheels go over a bump, the car will rotate around its transversal axis. The angle of rotation remains relatively small as the length of the car is its largest dimension; the higher weights like the engine bay are far from the centre of mass, resulting in a large inertial torque to counter outside forces. If all suspension elements of the wheels were connected hydraulically, the vehicle would absorb the bumps very efficiently (the rear struts compressed by the bump would deliver fluid into the front struts, resulting in immediate compensation: the rear would sink, the front would rise, restoring the horizontal position of the car). Unfortunately, this would also lead to slow transversal (dive and squat) oscillations, made even worse by acceleration, deceleration and varying distribution of weight inside the cabin.

As the inertia of the car body around its transversal axis is basically sufficient to counter the effect of longitudinal bumps, the front and rear suspension circuits are separated. The active height correction of the system acts as a further a non-linear stabilizer both countering dive and squat, and solving weight distribution problems.

On the other hand, if the bumps are transversal—for instance, a pothole under the right wheel and a bump under the left one—the car will rotate around its longitudinal axis. Being much less wide than long, the angle of rotation will be higher and the inertial torque is considerably lower to counter this kind of rotation. Completely independent sides would result in very little damping of roll movements: the low inertia provided by the body would find the reaction of the suspension too stiff. Hence, the two sides in the hydropneumatic suspension are interconnected, providing a push-pull operation of the two sides. The interconnection has special damping elements which react differently to different fluid movements between the sides: to quick suspension movements caused by potholes and bumps, or to slower changes occurring when driving in a curve.

To counter body roll resulting from the second, an additional element, an **anti-roll bar** is also needed. The effects of roll could be eliminated if the center of the roll could be identical to the center of the mass. As this is not possible, the opposite approach of moving the center of roll away from the center of mass could also help overcome body roll by increasing the opposing torque. This is the role of the anti-roll bar: similarly to a bike leaning into a curve, it lifts the inner side of the wheel, using the force on the outer edge, and this moves the center of roll outwards. In other words, the wheels and suspension elements do have roll, the role of the anti-roll bar is to isolate this roll from the body which should remain, ideally, horizontal. To accomplish this, the bar cannot be completely rigid (it has to absorb the road undulations without transferring them to the body), a torsion spring is the usual solution.

Such anti-roll bars are used on conventional spring suspension systems as well, however, there are substantial differences in the way the bar interacts with the rest of the sus-

pension on Citroëns. In a spring system, there is a considerable amount of interaction, a significant flow of energy in both directions between the suspension and the bar. The shock absorbers have to provide the damping for the anti-roll bar, introducing yet another interaction (in the hydraulic setup this is catered for by the damping inside the connection line between the sides).

Consequently, the hydropneumatic suspension has much less interdependence and compromise between damping, countering roll, squat and dive. In addition, it can provide solutions which are simply unfeasible mechanically in a conventional suspension. Cars with steel springs always have roll, including diagonal one, induced by undulations of the road—their anti-roll bar represent a constant mechanical connection between the sides, unable to differentiate between bumps and curves. Citroëns, on the other hand, have a varying interconnection depending on fluid movement—this is very easy to accomplish with hydraulics but extremely complicated with springs.

The only disadvantage is that damping occurs further from the source of the disturbance, and due to the good conductivity of sound via the hydraulic lines, this results in slightly more noise. The same effect makes the hydropneumatic suspension somewhat noisier than a conventional one. However, good sound insulation inside the cabin can help overcome this small annoyance.


This suspension layout reduces the **sensitivity** to under-inflated or blown tires and cross-wind. Even with largely uneven braking forces on the two sides the car will not pull to either side.

Although the hydropneumatic spring-absorber unit is an integrated unit from a technical point of view, hydraulics make it possible to place some hydraulic parts (for instance, the center spheres on Hydractive systems) in different locations, **reducing the amount of sprung mass**. Conventional springs have a considerable mass of their own while the mass of the nitrogen in the spheres is practically negligible. Even adding the mass of the fluid moving around in the system, the sum remains much below that of a steel spring. Hydropneumatic struts can be kept relatively small by increasing the operating pressure, which decreases the diameter of the struts. The automatic height correction reduces the mass further because the basic suspension mechanics can be simpler, without requiring multilinks and similar components.

The brakes share the mineral fluid with the suspension. This fluid boils at a very high temperature, therefore it provides great **resistance to vapor lock**. Due to the proportional regulation a hydropneumatic Citroën can keep braking as long as there is anything left of the brake pad. Even if the liquid starts to boil, there will be no vapor lock as the pressure is automatically released and remains proportional to the braking effort applied by the driver.

This system is often criticized for being overly complicated and prone to error, none of which accusations is true. The suspension is actually quite simple when considering its extra services in comparison to a conventional system and experience shows that the whole system is very reliable. The perfect functioning of the system relies mainly on the prescribed cleaning of the system and the change of the hydraulic fluid—adhering to these simple prescriptions can make the system very reliable.

A typical example: the BX




Finally, there are no forces in the suspension when the circuit is depressurised, allowing very **easy and safe servicing** of the relevant suspension and transmission parts.

Modern spring suspension systems are in fact capable of achieving some of these results. For instance, variable diameter or pitch springs coupled with hydraulic shock absorbers (incidentally, with a similar internal geometry as the damper elements used in Citroën spheres) behave similarly to these hydropneumatic units. The main difference is that even if these elements would be practically identical, all other functionality that comes either for free or at a small additional cost in Citroën systems—constant height, anti-dive, brake force regulation and so on—, require complex and expensive additional systems.

The illustration shows the basic layout of the suspension (differences on models fitted with power steering or ABS will be described in the corresponding chapters). Most components have an output line to collect leakage (which is intentional to keep the elements lubricated) and return it to the reservoir—although the outputs are indicated, the lines themselves are omitted for the sake of clarity. In reality, they are grouped together and go back to the reservoir.

The *high pressure supply subsystem* consists of a five-piston volumetric **high pressure pump** drawing the mineral suspension liquid called LHM from the **reservoir**. The fluid under pressure is stored in the **main accumulator**. It is the


task of a **pressure regulator**—built into the same unit with the accumulator—to admit fluid into the accumulator as soon as the pressure drops below the minimum value of 145 bar; as soon as the pressure reaches 170 bar, the regulator closes and the fluid continues its idle circulation from the pump, immediately back to the reservoir.



On simpler models the output marked with an asterisk is omitted and it goes to the return output inside the regulator unit instead, as shown by the dashed line. On models fitted with power assisted steering (DIRASS) this interconnecting line is missing and both outputs are used independently.

The spring below the piston ① is calibrated so that it will collapse only when pushed down with a pressure exceeding the cut-in threshold (145 bar). While the pressure in the main accumulator remains inferior, the piston stays in the upper position, allowing the pump to deliver fluid into the accumulator through the ball valve ⑤: the unit is *switched on*. The piston ② also remains in the upper position (its spring is calibrated to the cut-out pressure, 170 bar), letting the entering fluid fill up the chamber ③ as well. This, in turn, ensures that the piston ① stays in the upper position: the fluid pressure in this chamber plus the force of the spring counters the downward pressing force even if the pressure in the accumulator rises well above 145 bar.

The fluid supplied by the pump raises the pressure in the accumulator; as soon as it reaches 170 bar, its pressing force will exceed the retaining force of the spring under the piston ②, forcing it to the lower position. In this moment, the high pressure line coming from the another piston will be cut off and the fluid from the chamber ③ can escape back to the reservoir (yellow in the illustration).



BX

CX



With the back pressure now vanished from behind the piston ❶, the pressing force of the accumulator fluid drives it down at once: the regulator is *switched off* now. The fluid supplied by the pump returns back immediately: on PAS-equipped cars, to the flow distributor, on other vehicles, straight back to the LHM reservoir through the internal connection (dashed line).

Shortly, as the suspension and braking circuits start to use up the pressure in the main accumulator, the piston ❷ will return to its original position. Once there, the regulator is ready to start a new cycle.

The characteristic ticking which can be heard in Citroëns is the sound of the regulator pistons quickly moving one after the other, in quick succession: ❷ down, ❶ down, ❷ up. The opposite tick—❶ up, when the regulator is switched on to replenish the accumulator—is much softer.

The interconnection ❸ is normally closed. Opening it lets all the fluid stored under pressure return back to the LHM reservoir—this is the way the system is depressurized when any of the suspension elements need servicing.

The liquid—supplied to the rest of the system from the main accumulator—passes through a **security valve** whose task is to ensure safety by feeding the brake circuits first. The *front brake circuit* is always open but the other two outputs are blocked by a piston. If the pressure in the main circuit exceeds 100 bar, the fluid pushes the piston back against the force of the spring, opening up the suspension outputs as well. The electrical switch for the low hydraulic pressure warning lamp on the dashboard is built into this valve as well. This way, a sudden failure of the pump or the belt driving it will not leave the car without sufficient braking power.



The second circuit fed from the security valve is the *front suspension*. The fluid goes to the **front height corrector**. When the vehicle height is stabilized, the piston inside the corrector blocks the inlet of fluid, isolating the struts from the rest of the suspension. Body roll is limited by the damping effect of the restrictors built into the sphere supports and by forcing the fluid to run from the left to the right strut through a connection line. If the movement of the front anti-roll bar dictates that the front of the vehicle should be raised, the connecting linkage moves the piston upward, opening the inlet and letting additional fluid enter the front struts. When an opposite movement is required, the piston moves downward, letting the fluid at residual pressure flow back from the struts to the LHM reservoir. Both directions of flow are stopped and blocked when the height corrector piston resumes its middle position.

The mechanical connection between the anti-roll bar and the height corrector is not a rigid linkage but has some free play. Just before the height corrector, the connecting rod coming from the anti-roll bar hooks into a small window on the corrector side. Small movements of the control

rod do not change the position of the height corrector, only those are large enough to exceed this free play. In addition, the corrector has its internal (albeit low) resistance, besides, all rods are somewhat elastic, so in the end, all these factors make the height correction system filter out the higher frequency components of the suspension movement.


Observing an initial threshold which has to be crossed before any correction occurs not only reduces the strain and wear on the correctors but also prevents the system from developing self-oscillation. A powered system provides amplification and any feedback mechanism with a delay—such as the height correction—could potentially result in oscillations. The initial threshold ensures that there is no feedback, and consequently, no oscillation when the required correction is too small.

The next circuit is the *rear suspension*. Its layout and operation is identical to the front one, having its own height corrector.

The first circuit, as already mentioned, feeds the *front brakes*. The liquid under pressure flows into the **brake compensator valve**, operated by the brake pedal. In its neutral position, the brake circuits are connected to the return lines to ensure that the brakes are not under pressure. When the driver pushes on the pedal, this moves the first piston, closing the return output and opening up the outlet going to the **front brake cylinders**.

This piston and a spring behind it pushes the second piston which works similarly for the *rear brakes*, although those are not fed directly from the security valve but receive their supply from the rear suspension (later brake valves have three pistons but their method of operation is practically the same). In consequence, the braking force at the rear depends on the load: the more the back of the car is loaded, the stronger the rear brakes work. Actually, on a Citroën mostly used to carry only its driver, without much load in the trunk, the rear brake pads and disks wear much slower than those in the front.

The **damping elements** in the sphere supports consist of a central hole which is always open and additional small holes closed and opened by a spring as the flow of the hydraulic liquid dictates. Slower suspension movements like body roll, squat or dive result in a slower flow of the liquid and the smaller dynamic pressure differences are not sufficient to bend the spring cover open over the additional holes. The damping effect is therefore only determined by the diameter of the center hole.



The abrupt jolts caused by road irregularities, in contrast, cause faster flow. With the increasing pressure difference the fluid will open the spring cover and use the additional holes as well. This increased cross section results in a lower damping effect.

The additional holes are located in a circle around the center hole. There are two spring covers, one on each side, but they do not cover all the holes equally. Half of the holes (actually, every second one) are slightly enlarged on one side, the remaining half on the other side. By carefully adjusting the size of the holes, the designers could fine tune the damping factors independently for both directions of strut travel.

Hydractive I

The Hydractive I suspension system appeared with the XM. Unlike the simpler hydropneumatic suspension used on the DS, GS/GSA, CX, BX and some XMs, this one has two modes of operation, soft and hard. The suspension functions in soft mode but it will be switched to the hard mode when the computer deems this necessary for the sake of roadholding and safety.


XM

To achieve this, the first hydractive system adds two spheres (one for each axle) and an electric valve to the struts and spheres of the standard hydropneumatic setup.

During normal driving, the computer keeps the suspension in soft mode most of the time but—based on the input provided by many sensors (steering wheel, accelerator pedal, body movement, road speed and brake), including the *Sport/Comfort* switch on the dashboard—the suspension ECU decides when to switch to hard mode; in other words, when to deactivate the additional spheres for extra roadholding and safety.

When the driver selects the *Sport* setting, the suspension is switched to hard mode constantly. This setting is not what any Citroën driver would call comfortable... The successor system, Hydractive II overcomes this limitation.

The layout of the system (front suspension)



The illustration only depicts the differences to the standard hydropneumatic layout already presented in the previous section:

- ❶ A standard Citroën **sphere base** which fits a sphere without a damper block. The sphere volume and pressure differ for the front and rear, as well as according to the model of the car;
- ❷ A **hydraulically controlled isolation valve** that connects or isolates the sphere from the rest of the suspension, modifying the spring constant of the suspension;
- ❸ A **ball and piston valve** arrangement that limits fluid cross-flow between the left and right suspension struts in case of body roll. This valve is disabled for suspension height corrections, in order to guarantee that the fluid pressure in the corner struts remains equalized;
- ❹ Two **damping elements** similar to those used on the corner spheres, acting as dampers for the center one;


- ❺ An **electrically controlled valve** driven by the suspension ECU. In order to reduce heat build-up, the computer uses pulse width modulation to achieve a constant current through the coil. The initial voltage is higher to make the valve react quicker but it is reduced to a smaller value once the inductive effects have been overcome, should the valve stay on for a long enough time. The valve is capable of being on indefinitely when driven with this sustained current.

The front and rear suspension circuits are identical and the same electrovalve serves both subsystems.

Soft, hard, soft, hard...


The default electrical mode of the suspension, when the electro-valve ❺ is not energized, is *hard*.
???

SOFT MODE



While the computer keeps the suspension in **soft mode**, the electro-valve ⑤ is energized by the ECU, opening the feed pressure onto the isolation valve piston ② and by moving it, connecting the center sphere ① to the rest of the suspension. The fluid in the suspension has to pass through two damping elements ④ (one for each strut connection). When both struts move in unison, the center sphere behaves as a standard sphere with a damper hole twice as large as a single damper element, but when the car starts to roll, the fluid has to move from one strut to the other, passing through both damper elements consecutively. In addition to this double damping, the sphere ① itself acts as a damping string, absorbing quick changes in pressure between the two dampers. This dampens the body roll to some extent even in soft mode.

HARD MODE



Whenever the computer feels it necessary to switch to **hard mode**, it closes the electro-valve ⑤, not allowing the main feed pressure to move the isolation piston ②. The pressure inside the center sphere ①, always higher than that of the return path under normal operating conditions, will move the control piston into a position which closes off the center sphere completely. The remaining pressure in this sphere remains unknown but as the main circuit pressure might change while the suspension is in hard mode (due to either the dynamics of the suspension—acceleration, braking,

movement due to uneven surface—or the vehicle height altered by the driver), the computer equalizes the pressure periodically by enabling the control block to assume the soft position for a short period of time.

Hard mode serves three reasons. First, it provides higher resistance to body roll. The cross-flow of LHM from one strut to the other has to pass through both damper blocks as in soft mode, but it is additionally limited using the piston and ball valve ③, now switched into the hydraulic circuit between the damper elements instead of the center sphere. The ball is positioned in the fluid so that any cross-flow moves the ball and thus limits the flow, dampening the body roll as well.

Second, it limits dive and squat by helping out the height correctors. A stiffer suspension damps the vertical motion and therefore reduces the amount of correction required.

Third, hard mode not only limits the suspension travel between the body to the road but between the suspension elements and the body. Its aim is to reduce suspension movement at the cost of comfort but to gain safety, limiting the influence of the body movement to steering, very important in extreme situations like a flat tire.

When the vehicle is making a sharp left turn, tending to roll to the right, the right strut will be compressed and the left one expanded. The fluid is then forced from the compressed strut to the expanded one, moving the ball in the valve towards the outlet of the left strut; as soon as it reaches and covers the outlet orifice, it closes off any further cross-flow. The corner spheres are now isolated and has to provide all the damping themselves.




At the same time when the body roll is present, the car might need to change the ground clearance as well: for instance, when braking in a curve. The valve ③ therefore has an additional piston which lets the LHM flow between the circuits of the struts and of the height corrector. If the body has to be raised, the pressure in the height correctors will be higher than that in the suspension. This higher pressure pushes the piston, which in turn dislodges the ball and the pressure will raise equally in both struts (without dislodging the ball, only one of the struts would receive the fluid, resulting in incorrect operation).

If the body has to be lowered, the higher pressure in the struts will dislodge the ball again, opening the piston towards the return line and the fluid will escape from both struts, lowering the vehicle.

Sensory perceptions

The computer of the suspension system takes its input signals from the various sensors and based on a set of rules, dynamically activates the electric valve.

There are eleven inputs to the ECU. First, the **Comfort/Sport switch** on the dashboard, enabling the driver to choose between the two settings. The status light on the in-



strument panel informs about the setting selected (it *does not* indicate the mode the suspension is currently in).

The second input comes from a **vehicle speed sensor**. This inductive magnet tachogenerator generates 4 pulses per rotation, that is approximately 5 pulses per meter traveled (although this depends somewhat on tire size). It is located on the gearbox where the speedometer cable attaches, or in some versions, on the cable itself. The ECU determines the **acceleration** of the car by evaluating changes in vehicle speed for the duration of one second.

Another input arrives from the **steering wheel angle and speed sensor**, an optoelectronic device consisting of two infrared light beams, interrupted by a rotating disc with 28 holes. The ECU senses the quadrature signal changes of both sensors to effectively increase the resolution of the sensor (28 pulses per steering wheel revolution) by a factor of four. This produces one edge change every 3.214 degrees of steering wheel rotation. The direction of turning can be determined by the sequence of the edge changes.

To make decisions, the computer needs to know the straight ahead position of the steering wheel. The sensor does not have a built-in zero position (as it would not always work, due to misalignment and wear in the mechanical components). The computer uses heuristics instead:

First, the straight line position is assumed if the vehicle speed is above 30 km/h and the steering wheel position was not changed (an error margin of up to 4 pulses is allowed) for the last 90 seconds. Second, we know the maximum number of pulses in both directions from the center (lock to lock angle divided by two). If the steering wheel is found to turn more than this value (an error of up to 4 pulses is accepted here, too), this is a clear indication of an incorrect center reference: in this case the center position will be adjusted by the surplus.

The rotational speed of the steering wheel is determined by measuring the time elapsed between the individual pulse edges coming from the sensors.

A similar sensor informs the computer about the **movement of the car body**. Two infrared beams, the disc having 45 notches, similarly quadrupled by the ECU. Excessively long intervals are considered coming from slow height changes resulting from the driver selecting a different height setting, and are consequently discarded.

The sensor is connected to the front anti-roll bar, to the right of the height corrector linkage. Due to its location, it is capable of detecting both squat and dive, and to some extent, body roll. But as the sensor is mounted off-center, its sensitivity to roll is about three times less than the sensitivity to squat and dive. In all directions, it can measure both movement amplitude and speed of movement, using the same process as the steering wheel sensor does.

The **throttle pedal position sensor** is located below the dashboard, right next to the pedal mechanism, where the pedal can operate its sprung lever as it moves. The sensor is a potentiometer with an integrated serial resistor in the wiper's circuit.

The entire travel of the potentiometer is quantized into 256 steps by the analog-digital converter inside the ECU. The 5 V reference is supplied by the ECU itself. Due to the gas pedal initial position and maximum travel, about 160 to 220 steps out of 256 are being actually used.

The **brake pressure sensor** is a simple pressure activated switch located on a hydraulic conduit connector block, right next to the ABS block, at the bottom of the left front wing, in front of the wheelarch, under the battery. The switch makes contact at 35 bars of braking pressure.

The **door/tailgate open switches** are located on the door frame and in the boot latch. The door switches are all wired together in parallel and connected to one input line (and routed to the interior light dimmer and timer as well). The tailgate switch is connected to the other input line (and routed to the boot light and the tailgate opened detection input for the status display on the dashboard, too; the door open and bonnet open signals for the status display are generated by a separate set of switches, independent of the ones used for the suspension).

The usual **ignition switch** provides a power-on signal, triggering an internal reset and self diagnostic run in the ECU. Turning the ignition on and off also triggers internal events that guarantee proper pressure equalization between the center and corner spheres.

The brain behind the suspension

The ECU is a small microcomputer sensing the input signals coming from the various sensors. A very interesting and important aspect of the system is that it uses the driver of the car as a major part of its intelligence, making the operation very simple but effective. To achieve this, most of the sensors read the controls the driver operates.

The software contains the description of various conditions (status of the input lines and internal timers) governing when to activate-deactivate the electrovalve switching the suspension to either hard or soft mode. These conditions can be formulated as rules.

Every main input sensor has an associated rule: when the value collected from the sensor exceeds a specific threshold, the suspension is put into hard mode and the computer starts a timeout counter. For the suspension to return to soft mode at the end of the timeout period, the threshold must not be exceeded again during this time. If it was exceeded, the suspension stays in hard mode and the timeout starts all over again.

There are four additional rules overriding the normal operation—even if the sensor inputs call for a generic rule to be applied, these four conditions are checked first:

- ☒ the computer puts the suspension into soft mode when the ignition is turned on or off. This setting prevails until 30 seconds elapse or the vehicle speed exceeds 30 km/h, whichever comes first;
- ☒ if the computer determines any problem with its own operation or any of the input or output devices (including inconsistent values like no body movement but a vehicle speed above 30 km/h), the suspension will be switched to hard mode and stay there until the ignition is turned off or the doors are opened with the vehicle speed below 30 km/h. The ECU does run a self-diagnostic routine when the ignition is turned on but some sensors cannot be tested at this time, only during normal use;
- ☒ whenever the suspension stays in hard mode for more than one minute, the computer switches to soft mode momentarily to assure the equalization of pressures in the corner and center spheres. If the circumstances still

call for hard mode, the suspension will revert within 50 ms and restart the one-minute timeout period;

- below 30 km/h opening the doors or tailgate overrides any other rules and puts the suspension into soft mode to equalize the pressures in the spheres.

As already mentioned, the **steering wheel sensor** is used to derive two inputs values: steering wheel speed and angle. These values are treated separately with the purpose of calculating the lateral acceleration of the vehicle (vehicle speed, steering angle) and the potential change in this acceleration (vehicle speed, steering wheel speed). It is seemingly done this way to save memory which would otherwise be required for a full three-parameter lookup (based on vehicle speed, steering wheel angle, steering wheel speed). The steering wheel sensor rules actually give a measure of potential body roll. Body roll is significantly reduced in hard mode, consequently, the rules were set up to ensure that the body roll is minimized when there is potential for it, still the suspension stays soft to absorb bumps when there is no body roll caused by the vehicle changing direction.

If the **acceleration** or **deceleration** (braking) of the vehicle exceeds 0,3 g (approximately 3 m/s²) while the actual speed is above 30 km/h, the suspension will be switched to hard mode and a timeout of 1.2 seconds begin.

The table below shows the thresholds of **steering wheel angle** and **rotating speed**. If any of these values exceed the threshold for the actual vehicle speed, the suspension will switch to hard mode; it will revert to soft when the corresponding value drops below the threshold for at least 1 second if the switching was triggered by the steering wheel angle and 2 seconds if triggered by the rotational speed:

Vehicle speed (km/h)	Steering wheel angle (deg)	Vehicle speed (km/h)	Steering wheel speed (deg/s)
< 30	always soft	< 30	always soft
31–40	130	31–60	196
41–60	100	61–100	167
61–80	52	101–120	139
81–100	40	121 >	128
101–120	18		
121–140	15		
141 >	8		

The body movement amplitude and speed is derived from the output of the **body movement sensor**, although the two values are used in a different way.

The body movement speed is used as the parameter for the activation of two types of corrections:

- Flat tire correction:** if the body movement speed exceeds 300 mm/s, the suspension switches to hard mode, and all thresholds are modified to 60 mm. The timeout of the correction will be 0.4 s.
- Excessive body movement correction:** if the body movement exceeds 60 mm more than three times within three seconds, the suspension will switch to hard mode, and all thresholds are modified to 60 mm. The timeout of this correction will be 0.4 s.

The previous corrections stay enforced until one or more of the following conditions are satisfied:

- body movement amplitude remains under the modified threshold until the correction timeout elapses;
- suspension selector is set to the *Sport* setting;
- the vehicle accelerates above 159 km/h;
- the steering wheel angle exceeds the threshold value dependent on vehicle speed as specified in the following table.

Once any of these conditions are met, the suspension will revert to normal operation, with thresholds restored according to the table. Exceeding any of these thresholds will force the suspension into hard more. The computer checks every 0.8 seconds whether the conditions forcing the suspension into hard mode are still present, and if so, the system stays in hard mode.

Suspension down > 13 pulses, timeout 1 sec

Suspension up > 9 pulses, timeout 1 sec

Suspension change speed between 30 and 50 ms AND Durchfederung > 3 pulses, timeout 1 sec

Vehicle speed (km/h)	Dive (mm)	Squat (mm)	Steering wh pos (deg)	Vehicle speed (km/h)	Dive (mm)	Squat (mm)	Steering wh pos (deg)
< 30	—	—	—	< 30	—	—	—

The values delivered by the **throttle pedal sensor** are used with reference to the vehicle speed in order to anticipate the vehicle dynamics as a result of acceleration or deceleration. The rules for this sensor represent a reaction to probable vehicle squat (on acceleration) or dive (on deceleration). Both are significantly reduced when the suspension is in hard mode.

The suspension ECU quantizes the pedal position into five discrete steps: 0, 30, 40, 50 and 60 percent of the complete pedal travel. The computer measures the time elapsed as the pedal travels from one step to the next in either direction. If this time is inside the intervals shown in the table, the suspension will switch to hard mode. It will revert to soft if the pedal movement becomes slower for at least the duration of the timeout specified:

Pedal press speed (ms)	Timeout (s)	Pedal release speed (ms)	Timeout (s)
< 100	1	< 100	1
101–150	2	101–200	2

The **brake pressure sensor** detects the pressure in the front brake hydraulic circuit. Since this is a fixed threshold sensor, the suspension setting rule is simple: if the vehicle speed exceeds 30 km/h and the pressure is above 35 bar in the brake circuit, the suspension switches to hard mode. The system stays so to prevent excessive dive when brakes are applied while any of these two conditions are met (the timeout value is one second).

Without ignition and electrical feed to the suspension computer, the electro-valve would immediately return to hard mode. Loading or unloading the car, people getting in or out would induce pressure differences in the hydraulic system. These differences would equalize abruptly when the system is started again, causing the car to jump or sink vehemently. In order to avoid this, the computer allows an additional 30 seconds of timeout starting when any of the doors is opened or closed (as communicated by the **door**

and tailgate open sensors) , leaving the electro-valve energized for the duration of the timeout.

It is important to note that the suspension will switch to soft mode even with the ignition switch turned off. Early cars did not have this feature built directly into the computer but used an additional relay and circuits. On those models, the constantly energized electro-valve can drain the battery if the doors remain open for a long time. Starting with the H2 suspension computer (from ORGA 4860, February 28, 1990) the door sensors are observed by the ECU itself and the operation is enhanced with a 10-minute timeout period. After this interval, the electro-valves will always return to the hard, non-energized state.

Changing the state of the **ignition switch** provokes a transition to soft mode for a maximum of 30 seconds; reaching a vehicle speed of 30 km/h will cancel this mode prematurely. When the ignition is turned on, the ECU also runs a self-test diagnostic sequence lasting three seconds.

When the **suspension selector switch** is set to the *Sport* setting, all sensor inputs except for the vehicle speed sensor are ignored. Below 30 km/h the car stays in soft mode and switches to permanent hard mode above this speed. The **suspension status light** in the instrument panel has two functions:

- ⊠ when the ignition switch is turned on and the suspension set to *Comfort*, it will light up for the duration of the ECU self test. If the computer detects any malfunction in the course of this test, the light will flicker one or more times during this period;
- ⊠ when the suspension is set to *Sport*, the status light will remain lit to inform the driver of the setting chosen. The status light actually lights up or extinguishes only when the suspension rules have been changed in response to the mode select switch. This takes a short while because the internal timeouts are reset and some of the sensors are recalibrated. Because of this the light changing state is slightly delayed from the mode switch changing state.

Hydractive II

The second incarnation of the hydractive suspension appeared at February 1, 1993 (ORGA 5929). It was designed to overcome the biggest problem of the previous system, the very uncomfortable hard mode.

XM

Switching to *Sport* does not mean sticking to a hard, uncomfortable ride any more. On the Hydractive II, the relation between suspension modes and dashboard switch settings became more complicated: in both settings—*Normal* (the new name of *Comfort*) and *Sport*—the computer can switch to both hard and soft mode as it finds it necessary,


however, when set to *Sport*, the suspension becomes more sensitive and will sooner and more often switch to the hard mode.

Xan

1
2
3
4

Many models were also fitted with an anti-sink system that locks the system when the car is not running, using yet another sphere. Its only purpose is to keep the car from sinking when not used, it does not influence the functioning of the suspension system in any way.

The layout of the system (front suspension)



The center sphere circuits and supports were redesigned: they now house the electrovalves and the internal conduits serving the sphere were modified as well; the new control blocks connect, as previously, to the left and right corner spheres, the height corrector, and—depending on the control signal coming from the suspension computer—the center sphere. The elements are practically the same as on Hydractive I:



- ❶ A sphere base;
- ❷ A hydraulically controlled isolation valve;
- ❸ A ball and piston valve;
- ❹ Two damping elements;
- ❺ An electrically controlled valve driven by the suspension computer.

The front and rear suspension circuits are identical but hydraulically independent. The electro-valves are driven simultaneously, in parallel.

Trapped among pistons

The electro-valve ❺ is energized when the suspension is in its *soft* mode, hence, the default electrical position is *hard*. However, due to the indirect coupling between this valve and the isolation piston ❷ inside the control block, the hydraulics can stay in either position for extended periods of time with the electric valve disconnected, depending on the pressure differences between the strut and the main circuits. If the main suspension circuit has nominal pressure, the system stays in hard mode with the electric valve off or disconnected.

The two modes are practically the same as on the previous Hydractive system: in **soft mode** the electro-valve ❺ opens the feed pressure onto the isolation piston ❷ and by moving it, connects the center sphere ❶ to the rest of the suspension. In **hard mode**, the electro-valve ❺ closes and



lets the pressure inside the center sphere ① move the control piston into a position which closes off the center sphere completely.

The center sphere ① is now supplied directly from the height corrector in soft mode. This simplifies the ball valve arrangement with respect to Hydractive I.

Higher intelligence

The computer uses the same set of sensors as Hydractive I, the only difference is the **vehicle speed sensor** which is a Hall-effect sensor now. Its resolution have been doubled to 8 pulses generated per rotation, that is approximately 5 pulses per meter traveled (although this depends somewhat on tire size). It is located on the gearbox where the speedometer cable attaches, or in some versions, on the cable itself.

The internal algorithm of the computer became more sophisticated. While the Hydractive I had only one computer controlled mode (*Sport* switched the suspension to constant hard mode above 30 km/h of vehicle speed), the newer system has two such regimes of operation: in both *Normal* and *Sport* it dynamically activates the electro-valves of the suspension control blocks whenever it decides that the driving circumstances call for a firmer suspension. The difference is in the set of rules the computer uses to evaluate those circumstances: the rules are stricter for the *Sport* setting, with most of the thresholds reduced, thus the suspension will switch to hard mode much more readily.

The following table shows the thresholds of **steering wheel angle**. If the value observed by the sensor exceeds the threshold for the actual vehicle speed and the suspension setting, the suspension will switch to hard mode; it will revert to soft when the corresponding value drops below the threshold for at least 1.5 seconds:

Vehicle speed (km/h)	Steering wheel angle (deg)		Vehicle speed (km/h)	Steering wheel angle (deg)	
	Normal	Sport		Normal	Sport
< 34	—	—	90–99	33	22
34–39	174	119	100–119	26	27
40–49	100	67	120–139	23	15
50–59	84	56	140–158	20	13
60–68	68	45	159–179	13	9
69–78	55	37	179 >	10	7
79–89	42	28			

There is a similar table for the thresholds of the **steering wheel rotational speed** as well:

Vehicle speed (km/h)	Steering wheel speed (deg/s)		Vehicle speed (km/h)	Steering wheel speed (deg/s)	
	Normal	Sport		Normal	Sport
< 24	—	—	79–89	62	41
24–29	535	357	90–99	53	35
30–39	401	267	100–119	42	28
40–49	246	164	120–139	30	20
50–59	178	119	140–158	22	15
60–68	110	73	158 >	20	13
69–78	82	55			

The thresholds for **body movement** are:

Vehicle speed (km/h)	Dive (mm)	Squat (mm)	Steering wh pos (deg)	Vehicle speed (km/h)	Dive (mm)	Squat (mm)	Steering wh pos (deg)
< 10	—	—	—	100–109	48	48	13
10–33	84	60	—	110–119	48	42	13
34–39	84	60	87	120–129	48	42	11.5
40–49	54	48	50	130–139	42	42	11.5
50–59	54	48	42	140–149	42	42	10
60–68	54	48	34	150–158	42	36	10
69–78	54	48	27.5	159–179	42	36	6.5
79–89	54	48	21	179 >	36	36	5
90–99	48	48	16.5				

Note that the thresholds are the same for both *Normal* and *Sport* suspension settings

The thresholds of the **gas pedal sensor** are:

Vehicle speed (km/h)	Pedal press rate (steps/25 ms)		Vehicle speed (km/h)	Pedal release rate (steps/25 ms)	
	Normal	Sport		Normal	Sport
< 14	2	1.3	< 19	10	6.6
15-49	3	2	20-78	5	3.3
50-99	4	2.6	79-168	6	4
100-134	5	3.3	168 >	7	4.6
135-199	6	4			
199 >	7	4.6			

With the improved resolution of the vehicle speed sensor, the rules formerly referencing to 30 km/h are changed to 24 km/h. Thus, the suspension switches to hard mode if the **brake pressure sensor** detects a pressure above 30 bar and a vehicle speed in excess of 24 km/h.

Similarly, the suspension will switch to soft mode if the **ignition switch** is turned on, for a maximum of 30 seconds, but reaching a vehicle speed of 24 km/h will cancel this mode prematurely. It will switch to soft also if any **door** or the **tailgate** is opened but the vehicle speed is below 24 km/h. The reason for this is to equalize the pressure between all three spheres of an axle. Without it, the center sphere would retain its former pressure and once the vehicle exceeds the speed of 24 km/h, opening it would make the car jump or drop, depending on the actual pressure.

It is important to note that the suspension will switch to soft mode even with the ignition switch turned off. Should the doors remain open with the ignition switch in the off position, the suspension soft mode will be subjected to a 10-minute timeout period to avoid draining the battery as the soft mode requires the electric valves to be energized.

Anti-sink system

Many contemporary Citroëns—including both regular hydropneumatic and Hydractive Xantiae and XMs—have an anti-sink system (SC/MAC) fitted, to keep the car from lowering when not used. The system does not interfere with the normal functioning while in use. It attempts to minimize leaks inside the system by having only one element that can leak, the anti-sink valve itself.

XM

The introduction of this anti-sink valve coincided with the appearance 6+2 piston high pressure pump. As the suspension is fed from the smaller, two-piston side of the pump, pumping the car up from the low position would require a lot of time (although its performance is perfectly sufficient for the normal operation once the car is already running).

To avoid this scenario, the anti-sink valves fitted for each axle between the height corrector system and the suspension struts (or the hydraulic control block on Hydractive systems) keeps the car body from lowering when the engine is switched off. The valves operate on the pressure differences in the system, without any electrical control: when there is significant pressure in their control circuit, they keep their work circuit constantly open.

Under normal circumstances, the high pressure pump supplies the pressure regulator and the main accumulator with fluid. The output from these two feeds the whole system with high pressure, going through the security valve which keeps the brake circuit constantly under pressure, for obvious reasons of security. If there is enough pressure in the system, the security valve feeds the rest of the suspension via the anti-sink valves and the height correctors.



This pressure coming from the security valve appears in the control circuit of the anti-sink valves. When the car runs, the valves are constantly open, connecting the height correctors to the rest of the suspension and brake subsystems—everything functions exactly as in cars not equipped

with this anti-sink system. Even when the engine is turned off, the valves remain open as long as the feed from the accumulator remains at a higher pressure than that of the suspension. But as soon as the leakage in the struts, height regulators and the brake valve reduces the pressure in the main accumulator below the suspension pressure, the closing anti-sink valves isolate the suspension struts from the rest of the system. It is usually the front valve that closes first as the front of an unladen car is much heavier due to the engine and gearbox. Compared to a non-anti-sink car, the leakage is quite drastically reduced. For instance, a standard XM with its suspension in prime condition takes about 20-30 hours to sink completely, while with the anti-sink system this would take as much as ten days.

The rear anti-sink valve is connected slightly differently: in addition to feeding the rear suspension and the brake circuit, as usual, it connects to an additional anti-sink sphere as well. The function of this sphere is to maintain pressure in the braking circuit. As the brake valve is the most leaky element, it could exhaust the pressure between the piston and the plunger while the remaining pressure behind the piston (provided the high pressure and the front suspension circuits do not leak that much) stays rather high. In this case the anti-sink valve might open again in error—this additional sphere ensures that this will not happen.

This system maintains the car height by counteracting the internal leakage of the various suspension element that would make the pressure escape back to the reservoir. Elements that are in constant motion—height correctors, for instance—leak past their seals on purpose to lubricate themselves. The anti-sink valves—which move very rarely, need no intensive lubrication, thus are manufactured with very close tolerances and hardly leak themselves— isolate all the struts from the rest of the system to prevent any possible leakage to reduce the pressure in the struts, allowing the car to sink.

Xan



Activa Suspension

The Activa suspension—used only on some Xantia models—creates mixed feelings. Drivers requiring sporty handling and roadholding praise it because this car turns into curves without turning a hair: it stays completely horizontal and neutral. However, this comes at the expense of ride comfort.

The Activa system operates in two distinct steps. The first one is controlled mechanically by a **roll corrector** (the component is identical to the height correctors used in the suspension, see the details on page 26).

The corrector is connected by an L-shaped spring to the bottom wishbone. When the car takes a sharp left turn, its front left wheel will be forced down by the body roll caused by centrifugal force. As the wheel moves down, so does the end of its wishbone, pulling the linkage to the corrector. The piston inside the roll corrector moves upwards, opening the pressure feed into the stabilizing cylinders. These two cylinders are attached to the wheel suspension differently: in the front, the piston pushes the left wheel upwards while in the rear, the right wheel will be forced downwards. This diagonal correction counteracts the roll of the body.


Turning to the other side result in an inverse operation: the roll corrector opens the connection from the stabilizing cylinders back to the reservoir. The front left wheel moves downwards, the rear right one upwards, once again counteracting the effect of body roll.

An additional Activa sphere in the front acts as an extra accumulator but the rear sphere can be connected or decoupled electrically. Depending on the position of the piston inside the **electro-valve**, the high pressure feed is either allowed to reach the piston ② inside the control block, pushing it up and connecting the sphere ① to the rest of the circuit (dashed line on the illustration), or the residual pressure in the sphere moves the piston ② down, isolating the sphere ①.

When the Activa sphere is open to the rest of the system, roll correction is applied through a spring element formed by the accumulator and the Activa sphere. The supply side of the **stabilizing cylinder** pistons have half the area of the other side, connected to the Activa sphere ① with the valve ② open. Changes in the length of the linkage is therefore not transmitted directly to the roll bar. Upon the influence of external forces like body roll, the movement of the piston compresses the gas content in one sphere and at the same time, expands it in the other.

The stabilizing cylinder works as a spring with asymmetrical characteristics: its effective hardness is smaller around the corrected position, but it hardens progressively as the piston is forced out of that position.

The Activa system has two operating modes, depending on the position of the electro-valve ②. In the first mode roll correction is always active because the roll corrector is upset. The resulting flow of fluid will tend to move the active linkage upsetting the balance of pressure in the two extra




spheres, and making the coercive force be applied through a spring element which becomes progressively stiffer the more correction is needed.

The **ECU** controlling the electro-valve uses sensors identical to the Hydractive system. The values of vehicle speed, steering wheel rotation angle and speed determine when the second mode of anti-roll behavior has to be enforced. Similar to the operation of the suspension computer, the Activa ECU also uses the driver as the input to determine the motion of the vehicle body: if the roll is caused by the unevenness of the road surface, the steering wheel will not be rotated. In curves, the computer calculates the maximum potential lateral acceleration (vehicle speed is measured by its sensor, the turning radius is communicated by the steering wheel angle sensor, the mass of the car is a known constant—the centrifugal force can be calculated from these values) and decides whether the spring element formed by the two spheres needs to become rigid to make the system compensate for the body roll.

In this mode the Activa sphere is isolated from the rest of the system, the fluid line between the roll corrector and the active linkage is blocked at both ends, making the linkage completely rigid. Even if the roll collector end is open, the linkage remains quite rigid (providing for a very hard spring coupled with high damping); only half of the displacement escapes from the additional accumulator sphere through a restrictive regulator.

The additional damping of the Activa sphere is now switched off, the correction is applied only through the very hard roll-bar. When the possible range of correction is exhausted (strut linkage extends or contracts as far as it can), at about 0.6 g lateral acceleration, only the very hard roll-bar remains functional.


The diagrams showing the kinetic characteristics of an Activa car reveal the details. The first diagram shows the relationship between time and roll angle for a constant lateral acceleration. It can be observed clearly that the Hydractive system can only limit roll damping, not roll angle. Note that the initial slope of both Hydractive curves—the section up to 0.4–0.6 seconds—is practically the same in both soft and hard mode. This slope represents the combined hardness of the roll bar and the associated hydraulic components. Yet, the reaction time is longer in the soft mode (0.8 seconds versus 0.6, indicated by the last bend when the curve turns into a horizontal line). As the corner spheres are isolated and their combined gas volume is less in hard



mode, the maximum roll angle stabilizes around 2.5 degrees while in soft mode it reaches 3 degrees.

The second diagram depicts the relation between the lateral acceleration and the roll angle. The hydraulical-mechanical roll bar of the Activa starts the same as the Hydractive system with minimum lateral acceleration. But, while the Hydractive stays almost linear—the sharper you turn, the bigger the body roll angle will be—the Activa compensates by keeping the body roll angle at a constant below 0.5 degree up to a lateral acceleration of 0.6 g (by providing an effectively infinitely stiff roll bar setup). But even when the limits of the roll bar are reached, having contracted or extended it as far as it can go, the effective roll bar remains quite stiff: the roll angle will increase only moderately, up to a maximum of 1 degree.

...
...



Hydractive 3

The new C5 has a new suspension system, doing away with many solutions used on Citroëns for several decades, yet offering the same or even better comfort than before. Recent developments in electronics and computers made it possible to delegate many functions previously solved by mechanical-hydraulic components to electronic units.

C5

This third generation suspension system retains the same basic functioning as the previous systems. It also comes in two flavors: a simpler **Hydractive 3** reminiscent of the original hydropneumatic suspension of the DS-GS-BX-CX and a slightly more complicated **Hydractive 3+**, building upon the former Hydractive I and II (actually, Hydractive 3 is not hydractive in the sense we used this term before, its only special activity is to adjust the road clearance depending on speed and road condition).

Although the basic functioning is practically the same, the actual layout underwent significant changes. Most importantly, the previously mechanically operated height correctors became electronically controlled hydraulic units. And all hydraulic units except for the spheres—which were redesigned to give unlimited life expectancy—are now housed in a single unit, the **Built-in Hydroelectronic Interface (BHI)**. This compact unit has three main parts:

- the high pressure for the new synthetic fluid (called LDS, orange in color) is generated by a five-piston hydraulic pump **1**, driven by an electric motor (rotating at 2,300


rpm) operating independently of the engine, running only when necessary;

- the hydraulic units, including an accumulator **2** to even out the pressure pulsations of the pump, four electrovalves **3** and **4** and two hydraulic valves **8** serving the height regulation and anti-sink behavior, some in-line filters **5** and an overpressure valve **7** (taking the role of the pressure regulator of previous systems).

- the electronic computer **6**, communicating with other computers across the multiplex network to read the inputs of various sensors and to control both the HP pump motor and the electrovalves.

In contrast to the height correctors of previous systems, operated mechanically via a linkage coupled to the anti-roll bars, the new system used electronic sensors to learn the actual height of the suspension and electric actuators to modify the ground clearance whenever needed. The main advantage of using them is that the ECU can implement very sophisticated algorithms to derive and apply height correction, what were impossible with the mechanically linked feedback with simple thresholds.

The computer **6** is connected to the CAN multiplex network, providing access to the messages sent by the BSI and its fellow computers controlling the engine and the ABS. The inputs the suspension ECU uses comprise of rear and front body height, brake pedal, vehicle speed and acceleration, open-closed status of the doors (including the tailgate), plus the steering wheel angle and rotating speed on the Hydractive 3+.



As usual with Citroëns, the driver can select from four height settings (although the selector is no longer mechanically coupled with the hydraulics, it is a simple electronic switch sending signals to the computer): *high*, *track* (plus 40 mm), *normal* and *low*. The selected setting is displayed on the multifunction screen in the dashboard. The computer also prevents unsuitable settings being selected. The *high* option is not available when the car is traveling faster than 10 km/h and neither *track* nor *low* mode can be selected above 40 km/h.



In addition to the manual settings, the system adjusts the ground clearance automatically. Below 110 km/h on well surfaced roads the ride height remains standard but as soon as this speed is exceeded, the vehicle will be lowered by 15 mm at the front and 11 mm at the rear. This change lowers the center of gravity, improving stability, lowering fuel consumption (by reducing drag) and reducing the sensibility to crosswinds. The car resumes the standard ride height when its speed drops below 90 km/h.

On poorly surfaced roads (the computer learns about the road quality by monitoring data on vehicle speed, height and movement of the suspension) the ride height will be increased. The maximum increase would be 20 mm but this setting is only used on very poor roads and with the vehicle traveling below 60 km/h.

The general height of the vehicle (filtering out rapid movements due to suspension travel) is checked, and if necessary, adjusted every 10 seconds and when any of the doors is opened or closed (even with the ignition switched off).

Hydractive 3+

Just like its predecessor, this system also has two modes, *firm* and *soft*. A **stiffness regulator**—an additional sphere and a hydraulic control block per axle—isolates or connects the corner and center spheres. Its functioning is practically equivalent to the similar control block of the Hydractive II: the computer controlled electro-valve ④




opens the feed pressure onto the isolation piston ② and by moving it, connects the center sphere ① to the rest of the suspension, switching the suspension to **soft mode** of the suspension. Closing the electro-valve ④ obstructs the hydraulic supply coming from the BHI; the residual pressure in the center sphere ① moves the isolation piston ② downwards into a position which closes off the center sphere completely: the suspension switches to **hard mode**.

The suspension has two settings the driver can choose from, *Normal* and *Sport*. The new stiffness regulators together with the center spheres are isolated in hard mode and re-activated in soft mode in response to the various inputs received and processed by the suspension ECU. The functioning of the computer is basically similar to the Hydractive II ECU: it uses tables and rules to set up thresholds on the value on many sensor inputs to determine when to switch to hard mode. Just like on its predecessor, the *Sport* setting does not mean constant hard mode, just lowered, more sensitive thresholds for the switching.

The computer observes the following input parameters: the **height** and **sport settings** specified by the driver (communicated by the BSI); the **vehicle speed** and the **longitudinal-lateral acceleration** of the body (communicated on the CAN), the **angle** and **speed of rotation of the steering wheel** (the type of the sensor depends on whether the car is equipped with ESP, in this case the sensor connects to the multiplex network instead of directly to the suspension ECU), the speed of **suspension travel** (using the values of the front and rear height sensors), the open-closed status of the **doors** (communicated by the BSI) and the movement of the **accelerator pedal or butterfly**.

Steering



Power Assisted Steering

The PAS steering (DIRASS, Direction Assistée) used on Citroëns is not radically different from similar systems on other cars. Naturally, having a high pressure hydraulic system at disposal influences the layout.

BX

The fluid requirements of the various hydraulics subsystems differ significantly: while the brakes require only a very little amount of LHM and the suspension somewhat more, the power steering cannot work without large amounts of mineral fluid provided at a moment's notice. A **flow distributor** built into the first hydraulic circuit—that of the hydraulic pump, the main accumulator and the pressure regulator—controls the hydraulic pressure between the steering circuit and the suspension-brake circuits on PAS cars.

The rest is rather simple. A hydraulic **ram cylinder** is mounted on the rack of a traditional rack-and-pinion steering gear unit. The pressure of the hydraulic fluid supplied to assist the driver in turning the steering wheel is controlled by the flow distributor and a **control valve**. The flow distributor has the following components:

- ❶ a **slide valve** to divide the amount of fluid;
- ❷ another **slide valve** to limit the amount of fluid;
- ❸ a **pressure limiting valve** to limit the pressure of the LHM when the steering wheel is turned completely to lock;


The steering control valve has three important elements:

- ❹ a **distributor** mounted to the pinion;
- ❺ a **rotor** fixed on the end of the steering rack;
- ❻ a **torsion bar** between the distributor and the rotor.

On the main illustration, the power assisted steering system is shown when it operates with the steering wheel in the straight-ahead position and the **pressure regulator** is switched on. The slide valve ❶ inside the flow distributor divides the mineral fluid coming from the high pressure pump between the main and the steering hydraulic circuits (the main circuit having priority). Both the distributor ❹ and the rotor ❺ are in neutral position—the torsion bar between the two is not functioning). Both chambers of the ram cylinder are fed without pressure. All the fluid arriving through the distributor flows back to the LHM reservoir.

When the pressure regulator switches off while the steering wheel still is in its straight-ahead position, the pressure starts to rise until it reaches 170 bar again and disconnects the feed to the **main accumulator**. The main slide valve of the pressure regulator is connected to the second feeding channel of the flow distributor. All the fluid supplied by the HP pump now feeds the flow distributor where the slide valve ❷ is responsible for limiting the amount of fluid transported by the control valve. The whole amount of fluid still returns to the reservoir.


Now let's assume the driver starts to steer to the right. The rotor ❺ starts to rotate with reference to the distributor ❹. The control valve closes the path of the fluid coming from the flow distributor which no longer is allowed to enter the valve. The pressure begins to rise in the circuit between the control valve and the flow distributor, moving the slide valve ❶, which in turn modifies the ratio of fluid, favoring the PAS circuit. The fluid will enter the right chamber of the ram cylinder while the left chamber can be emptied



into the reservoir via the rotor of the control valve. This pressure difference moves the piston 7 to the left inside the cylinder, helping the car to make a right turn. If the steering wheel stays at the right lock, the pressure limiting valve 3 inside the flow distributor maintains a maximum pressure of 140 bar—when the pressure rises above this value, the fluid pushes the ball of the valve backwards, sending the excess fluid back to the reservoir.

When the steering wheel is turned to the left, the rotor 5 rotates in the opposite direction. It starts by cutting off the return of fluid to the main reservoir. The pressure will rise again in the circuit between the flow distributor and the control valve. The rotor allows the LHM to enter both chambers of the steering ram, however, the pressed area of the left chamber is twice as large as that of the right chamber, thus the piston will move to the right, helping the car turn to the left.

The hydraulic assistance is only needed while the driver is actually turning the steering wheel. When the rotating force on the steering wheel ceases—the driver has finished turning the wheel—the angle difference between the distributor 4 and the rotor 5, made possible by the flexibility



of the torsion bar 6 disappears, reverting the system to the neutral position, stopping the power assistance. When the driver releases the steering wheel back to the straight-ahead position, an opposite operation will start.

Later XMs and Xantiae omit this distributor and use a two-section high pressure pump with two independent out-

XM

DIRAVI Steering

Another gem of engineering, the DIRAVI steering, made its debut on the SM, excelled in many CXs and the flagship, V6 XMs (left hand only, the small amount sold in the UK never justified the expenses of the conversion to RHD).


The DIRAVI (Direction Rappel Asservi, Steering with Limiting Counterforce) steering is as unique as the hydropneumatic suspension—it was never used by any other manufacturer, although its excellence over conventional power assisted systems speaks for itself.

As usual, it has some quirks confusing the average driver during their first meeting. First of all, it is geared very high: it only took two turns of the steering wheel from lock to lock (one turn for each side) to steer on the SM. Later models, the CX and the XM retained this feature although the number of turns was larger (2.5 and 3.3). The gear ratio could have been much higher, the engineers themselves insisted on a single turn lock to lock for the SM (which would, interestingly, void the need for a circular steering wheel completely). The final solution was a compromise to reduce the initial strangeness of the steering for the drivers already accustomed to traditional systems.

Certainly, making the gearing so high is not complicated in itself but a conventional (even power assisted) system with such rapid a response would be unusable. As the car obviously has to have a similar turning circle as other cars, too responsive a steering would mean that even the slightest movement of the steering wheel would induce excessive deviation of the car from the straight line. To avoid this, it

uses an opposing force, increasing with the vehicle speed. With this setup, in spite of the very high gearing, it is very easy to use it during parking, yet it offers exceptional stability at high speeds: it actually runs like a train on its rails, requiring a sensible amount of force on the steering wheel to deviate it from the straight line. And an additional feature: the steering wheel (and the roadwheels, naturally) center themselves even if the car is stationary.

Second, there is no feeling of feedback from the road through the steering wheel. Other steering systems have a constant mechanical connection between the steering wheel and the roadwheels, the DIRASS only adds some force to the one exerted by the driver. DIRAVI is different: simply put, the usual path between the steering wheel and the steering rack is divided into two halves, with a hydraulic unit in the middle. When the driver turns the steering wheel, this only operates the gears and valves in the hydraulic unit. The hydraulic pressure then moves the steering cylinder and the roadwheels. The lower half of the mechanics works in the opposite direction, as a negative feedback, returning the hydraulic system to the neutral position as soon as the wheels reach the required direction. The hydraulic cylinder and the wheels become locked, no bump or pothole can deviate them from their determined direction. Note that this neutral position is not always the straight-ahead direction, the hydraulics return to neutral whenever the steering wheel is held at a given angle for any longer period of time. Letting the steering wheel rotate back or turning it further in the previous direction will initiate a new mechanical-hydraulic cycle as described above.




Thus, the lower mechanical link, the feedback from the roadwheels does not extend beyond the hydraulic unit. Everything the driver feels is generated artificially. One drawback for uninitiated drivers is the lack of noticeable feedback indicating that the wheels are skidding or driving in a ditch. The driver has to learn to feel the behavior of the car via other sensory means and this is probably the main reason why anyone not prepared for a period of learning will immediately dislike DIRAVI. But once accustomed to the system, it is more ergonomic and stress-free than any other steering system.

The DIRAVI system uses four main components:

The **steering rack and hydraulic ram cylinder** with a piston inside. The areas on which the pressure acts on the left and right sides of the piston are different—the left one is twice as large as the right one—, thus to keep the piston in neutral position, the right hand side must have twice as much hydraulic pressure than the left hand side. As this side is fed from the high pressure of the hydraulic system, a control unit manipulates the pressure on the other side.

This **steering control unit** is connected to the steering column. It has a coupling ④ inside which is very loosely connected, with a significant amount of free play (nearly 30 degrees). Under normal circumstances this coupling stays in the middle, so the free play is irrelevant but it serves as a mechanical backup for safety if there would be any failure in the hydraulic system. In this case, the car can be steered mechanically, although much heavier and with a large free play on the steering wheel.



The main illustration shows the steering system with the steering wheel in the straight-ahead position. When the driver rotates the steering wheel, the steering column turns the gear ① inside the control unit. The set of levers ② attached to this wheel transform the relative rotation (relative to the previous hydraulically stabilized steering wheel position) of the steering wheel into a horizontal motion: turning the steering wheel to the left pulls the slide valve ③, letting the high pressure fluid enter the left chamber of the cylinder. The right chamber is constantly at this same pressure, however, the area on the left side of the piston is twice as large as on the other side, thus the resulting higher force will move the steering rack to the right, turning the roadwheels to the left.

If the driver rotates the steering wheel to the right, the levers ② push the slide valve ③, draining the LHM from the left chamber of the cylinder back to the reservoir. As the right chamber is still under the constant pressure, the resulting force moves the rack to the left, thus the car starts to turn to the right.

As we have already mentioned, the moving steering rack rotates the pinion and—through the steering feedback—the cogwheel ②. The levers linking this gear to the valve ③ now work in the opposite direction, returning the valve to its neutral position, cutting off the LHM supply to the steering rack. The roadwheels stay in the angled position corre-

sponding to the position of the steering wheel; due to the closed valve ③, the steering gear and the roadwheels are hydraulically locked, resulting in high turning stability.

To make the steering progressively heavier as the speed of the vehicle increases, the **steering centering pressure regulator**—a centrifugal device—is driven by a cable from the gearbox. Its spinning weights open up a slide valve ④ admitting some fluid from the high pressure circuit into the centering device, or closes it to drain the extra fluid back to the reservoir.

The faster the car runs, the bigger is this hydraulic pressure sent to the **steering wheel centering device**. This consists of an eccentric cam ⑤ geared to the steering wheel side of the unit, with a ratio making it turn less than a full turn while the steering wheel is rotated from lock to lock. A piston ⑥ forced down by the mentioned hydraulic pressure pushes a roller ⑦ against this cam. Being eccentric, the only stable position is when the cam is centered. The centering force can be regulated by changing the hydraulic pressure behind the piston.

The hydraulic pressure behind the piston ⑥—being dependent on the vehicle speed—represents the progressive counter-force needed to make the steering gradually heavier at highway speeds. In addition, it returns the steering gear to the neutral, straight-ahead position when the driver releases the steering wheel. While the wheels of a DIRASS car return to the center themselves, forcing the rack and steering wheel as well, on DIRAVI the opposite is true: the force on the angled wheels is attenuated infinitely, having no influence whatsoever on the steering wheel. This additional device returns the steering wheel to the center instead, just as if you have turned it back yourself.

During the rotation of the steering wheel, the lower piston was pushed up by the roller ⑦ and the eccentric cam ⑤. The fluid leaves the chamber through the ball valve now opened. While this piston moves upwards, it compresses the spring, which in turn pushes the upper piston slightly up, freeing the calibrated bore ⑩.

As soon as the driver releases the steering wheel, the opposite of the previous operation takes place. The ball valve will be closed by the entering fluid, thus the LHM has to go through the center bore of the upper piston, leaving via the calibrated bore ⑩. Due to this resistance, it carries the upper piston down slightly, compressing the spring. This downward force pushes the lower piston together with the roller ⑦ down, and the torque exerted on the eccentric cam ⑤ forces that to rotate back into its neutral position, returning the complete steering gear to the straight-ahead position. At the end, the spring will return the upper piston to its original position inside the centering device. The restriction of the bore ⑩ keeps the steering wheel from returning to the center position too fast.

The last component is an **adjustment cam** allowing the adjustment of the pinion relative to the disk on the pinion end of the steering column.



Self-steering Rear

sasasa

Brakes




Standard braking system


The brake compensator valve is the most complicated part of the hydropneumatic suspension. This component is not only responsible for the operation of the brakes but forms a central part of the whole system.

BX

With no pressure applied on the pedal, the brake valve connects both the front and the rear brake circuits to the operational return, ensuring that the brake pads retract from the disks. Note that both braking pistons (4 and 5) have internal conduits (drawn as Y-shaped lines) offering a passage for the brake pressure to escape into the return line:




At the same time, the pressure coming from the rear suspension exerts pressure on the right hand side of the piston 1. As soon as this pressure exceeds a given (albeit quite low) level, the piston will be forced to the left. The cup 2 will be forced to the left as well which, in turn, will compress the spring 3. As soon as the piston 1 starts to move, the lip on its face will move away from the edge of the housing, effectively increasing the area the hydraulic liquid can exert its pressure upon. This setup introduces some hysteresis in order to prevent unwanted oscillations in the system: although the piston 1 and the cup 2 moves as soon as the pressure exceeds the threshold level, it stays in that position even if the pressure drops a little:




The reason for this is to lower the counter force on piston 4, forming the valve for the rear brakes. This makes the

rear brakes extremely sensitive at the very beginning of braking, minimize their delay in order to allow the anti-dive behavior of the rear trailing arm suspension to take effect.

As soon as the driver starts to brake by pressing on the brake pedal, which in turn presses the front piston 5 behind the rubber boot at the end of the unit, the feed from the main accumulator becomes connected to the front brake circuit:



In addition, the high pressure liquid will also pass through the channel in the middle of the front piston 5 and enters the chamber between the front 5 and rear 4 pistons. The pressure of the LHM will push the rear piston 4 to the left, too, which in turn activates the rear brakes:



The presence of pressure in the rear suspension has kept the piston 1 and cup 2 pushed back to the left (against the force of the spring 3) until now. But the increasing pressure in the rear brake circuit, being present on the left side of the piston 1 as well, will exert a counter-force, forcing the piston 1 to the right. Via the passage in the rear piston 4, the pressurized LHM of the rear suspension enters the chamber between that and the piston 1 (all these pistons have lipped faces to prevent unwanted oscillations by introducing hysteresis action, as described above).

We have discussed piston 1 as a single component so far but in reality, it is made of two parts. We introduce part 6 here so that we can refer to its right side (forming a chamber with the left side of the rear piston 4). In the following

formulae, piston ① will only denote the part in direct touch with the pressure coming from the rear suspension. Thus, the force on the left side of the piston ① comes from both the spring ③ and the braking pressure (blue) exerted on the left side of piston ①:

$$F_{\text{left}} = F_{\text{③}} + p_{\text{brake}} A_{\text{①left}}$$

As we can see from the cross section of the valve, the combined areas of ⑥ and ① right are equal to that of ① left:

$$F_{\text{left}} = F_{\text{③}} + p_{\text{brake}} (A_{\text{①right}} + A_{\text{⑥right}})$$

Similarly, the force on the right comes from the suspension pressure (green) on the right side of piston ① as well as the braking pressure (blue) on the right side of piston ⑥ (arriving via the channel in the rear piston ④):

$$F_{\text{right}} = p_{\text{rear}} A_{\text{①right}} + p_{\text{brake}} A_{\text{⑥right}}$$

Brake force regulation is performed by the movement of pistons ①⑥ until an equilibrium of forces is obtained on their two sides, and the consequent influence on pistons ④ and ⑤. Limiting of rear brake pressure with respect to rear suspension pressure is insured by the force generated from the pressure in the rear brake circuit being assisted by the spring ③. Because of this spring, that pressure is always smaller by an amount proportional to the spring tension:


$$F_{\text{left}} = F_{\text{right}}$$

$$F_{\text{③}} + p_{\text{brake}} A_{\text{①right}} + p_{\text{brake}} A_{\text{⑥right}} = p_{\text{rear}} A_{\text{①right}} + p_{\text{brake}} A_{\text{⑥right}}$$

$$p_{\text{brake}} A_{\text{①right}} = p_{\text{rear}} A_{\text{①right}} - F_{\text{③}}$$

$$p_{\text{brake}} = p_{\text{rear}} - \frac{F_{\text{③}}}{A_{\text{①right}}}$$

The balance of forces depend on the spring, as well as the suspension and brake pressures. Limiting occurs whenever the rise in rear brake pressure causes the pistons ①⑥ move to the right:



As the LHM trapped between the rear ④ and front ⑤ pistons is not compressible, the second one will also move to the right, providing a counter-force to the pressure of the driver's foot. As this movement will partially close the passage between the main pressure feed and the front brakes, the pressure in the front circuit is also forced to reach an equilibrium between the force applied to the brake pedal and that of pushing the piston ⑤ to the right.

The springs between the pistons ⑥, ④ and ⑤ are not taking part in this process. They are small, soft springs, aiming


only to center the rear piston ④ between the other two but their effect is otherwise negligible.

From this description, we can draw a few conclusions. First of all, the maximum rear brake pressure can never exceed that of the rear suspension. Consequently, if there is no rear suspension pressure, the rear wheels do not brake.


When the driver starts to brake, the rear brakes will bite immediately as the suspension pressure overpowers the spring ③, pushing the pistons ①⑥ to the left and allowing the rear piston ④ to open quickly and completely. This piston will also balance the forces in the front and rear brake circuit, so, during this initial braking period, the rear brakes operate just like the front ones. As the rear brake pressure builds up, it will always be dynamically limited with respect to the rear suspension pressure, as described by the formula above.

Stop breaking, please...

The brake compensator valve used on early CXs was slightly different from the one described above, without the built-in brake force limiting effect of the piston ① and the cup ②. These CX Break models had a separate rear **brake force limiter** to provide this functionality.



When there is no pressure in the rear suspension (for instance, the suspension is set to low), the force of the spring ④ keeps the piston in the neutral position, completely closing the feed to the rear brakes from the brake compensator valve.




When the suspension is under normal pressure, the force ① supplied by the rear suspension fluid exceeds the counter force ② provided by the spring. The piston stays in the open position, letting the fluid pass to the rear brakes. As soon as the driver starts braking, the force ② increases by the additional pressure coming from the front brakes, via the LHM entering through the ball valve ③.

As soon as the incoming front brake pressure exceeds the rear suspension pressure by more than 28 bar (in other words, the combined force of front pressure and that of the calibrated spring ④ becomes larger than the rear suspension pressure), the piston moves again to the left, cutting out the additional pressure to the rear brakes, which will then continue to brake with this constant pressure.

To avoid a sudden cut-off of pressure, the ball valve ③ is combined with a damping element ⑤ to smoothen the changes. This element is a coil made of hydraulic piping, closed at the far end. The elasticity of the pipe and the coil acts as a small hydraulic spring capable of lowering the resonance of the chamber, hindering unwanted oscillations.

Anti-lock Braking System

Models with higher performance level came fitted with ABS.



The principle of operation is the same as on cars with conventional braking systems but the layout is much simpler as all we need to control the operating pressure of the brakes are a few electro-valves.

BX

During braking, the **ABS computer** monitors the changes in the rotational speed of each roadwheel, communicated by inductive magnetic **sensors** reading the individual cogs of a toothed wheel fitted inside the cavity of the brake discs. The computer does not interfere with the braking if the vehicle speed (as measured with the same sensors) is below 5 km/h.

If any of the wheels begins to slow at a faster rate than the others, the ABS reduces the hydraulic pressure fed to the brake caliper of the wheel in question to avoid the wheel being locked. Although every wheel has its own sensor, the rear brake calipers receive the same pressure, only the front ones are fed separately. As soon as road grip is re-

gained, the hydraulic pressure to the brake will be restored. The computer is capable of cycling the pressure with a frequency of several times a second.

To actually control the pressure, the system uses a three-unit hydraulic block (one block each for the front brakes, one for both rear brakes). All three units comprise two electro-valves, an inlet **1** and a return **2** valve.

During the rising period of normal braking, without the need for the intervention of the ABS computer, the brakes operate in **phase 1**: the inlet valve **1** is open but the return valve **2** is closed. The braking functions as in a system without ABS: the incoming hydraulic pressure is directly routed to the brake caliper.

Under constant braking (**phase 2**) both valves close to maintain a steady hydraulic pressure in the brake calipers.

When the ECU senses the need for intervention, the electro-valves proceed to **phase 3**: the inlet valve **1** closes while the return valve **2** closes. Hydraulic pressure will be released from the brake caliper, reducing the braking force. To restore the braking effort, the ECU will return to phase 1 in a short while.


The ABS computer has a built-in diagnostic feature, checking the components both when the ignition is turned on and during braking. Any failure will be reported by a warning lamp or a warning message of the board computer. As you can see from the illustration, the springs inside the valves are located in such a way that the mechanical default mode is phase 1—the normal braking—for all three hydraulic blocks. Any failure in the ABS system will therefore revert it to the usual, non-assisted braking.

CX


Early CXs has a slightly different ABS system. The general layout is the same, but the hydraulic block only has three valves, one for each brake circuit, however, they have three positions. Without energizing current, they route the fluid coming from the brake accumulator to the brakes. In phase 2, a medium current switches them to isolate the brake calipers, while a larger current opens it completely to let the pressure escape from the brakes into the return lines.

XM

On XM the hydraulic block has five electro-valves only. I am not sure how they connect internally, but I suspect that the valve that closes supply for the front brakes is common for the left and right wheel ???



Electrical Systems



Multiplex network

Circuit layouts already universally adopted in computers finally made their way into contemporary cars. Although their functioning might be frighteningly complex for people used to traditional circuits, they actually make the cabling very simple and the addition of component interactions possible in ways never experienced before.

Conventionally, cars used individual wires connecting the various elements—steadily increasing in number—on board. The huge amount of wires, connectors, wiring harnesses were a constant source of connection problems. The various circuits were largely independent (sharing only the feed and the ground), although some components had to interact (for instance, fog lights should work only when the headlights are switched on), necessitating connections between the various components (usually using some kind of a switching logic, relays for simpler tasks and small electronic modules for more complicated ones).

As various subsystems (engine management, suspension, ABS, etc.) came from different manufacturers, some functions were even built in parallel. Several subsystems might rely on the signal sent by a coolant temperature or a vehicle speed sensor but it was simpler for the manufacturers to fit two or three such sensors into various places, using every one of them only by their respective subsystem, than to find ways to share the sensors, introducing interconnecting wires and the danger of one failing subsystem to influence the others.

The multiplex wiring first seen on late XMs and later used on newer models like the Xsara Picasso or the C5 introduces a radically different concept: just like in the computer used to read this book, there is a central backbone circuit called *bus* which goes around the whole car—actually, there are four of them, a **Controller Area Network (CAN)** and three **Vehicle Area Networks (VANs)**, dealing with different areas: the CAN is only responsible for the connection between the central unit and the engine, gearbox and suspension computers, the VANs for the rest of the systems: the first serves the safety systems like the airbag, the second the various doors (including the sunroof) and the anti-theft system, the third everything else: the instrumentation and the comfort gadgets.


The bus—in contrast to the traditional wiring harnesses hosting many individual wires running side by side to serve different components—is a common channel of information flow for all components connecting to it. It uses only two wires which all associated components connect to in parallel (in addition to this, the devices are connected to the ground as usual; the two input wires serve as a safety measure, using them both makes the system resistant to any outside interference, and the whole system remains functional even if one of the bus wires becomes broken, shorted to ground or positive feed). There is no special controller or owner of this bus, each device connecting to it is free to send or receive messages and commands to the others, at a rather high speed (approximately ??? messages per second).

XM

C5

Xs

Buses in the C5



Each **message** or command is a sequence of a few numbers, specifying:

- ✦ the sender and the intended recipient of the message (every device connecting to the multiplex bus has its own address, a unique numerical identifier—for instance, the fuel level sensor has the address 4315, the instrument panel is 0004);
- ✦ whether the recipient should acknowledge the message as it processes it;
- ✦ the actual data the message transmits;
- ✦ some additional values to check the integrity and validity of the message at the receiving end.

Each major unit sends its own data into the network at pre-determined intervals, marking the message with its own address as a sender (some simpler sensors are connected directly to a computer which sends the messages relating to their measured values on their behalf). With our example, the fuel level sensor sends the amount of fuel it measures, specifying the central unit (BSI) as the intended recipient. As soon as the BSI sees this message circulating on the network, it processes it by retrieving the data—the value of fuel level—from the message and comparing it to the previously known value. As the amount of fuel is not supposed to change drastically from one moment to the other, it discards the new value if it differs too much from the previous one.

If the new value is acceptable, the BSI emits another message of its own, addressed to the instrument panel this time. As the instrument panel receives this second message, it extracts the data representing the amount of fuel left in the tank and turns this signal into the physical rotation of the gauge needle.

All devices are constantly observing the bus for messages addressed to them, ignoring the ones sent to other recipients (although there are special broadcast messages sent to all devices, without specifying a single addressee)—actually, the instrument panel saw the original message coming from the level sensor as well but ignored it, it only acted when the second message, sent by the BSI and addressed specifically to it, arrived.

All components work in a similar way. Some are simple enough to send a few simple messages (like sensors or switches) or to receive only a few ones (like electric window motors). Others are complex subsystems themselves, like the suspension, observing the input from a large number of sensors and performing complex operations. But as they are all connected to a common bus, the possibility of interaction is already there. Whether the headlights light up, the electric windows close and the wiper starts to work in case of rain, or whether the passenger side external rear view mirror folds down when engaging reverse gear have all become simple questions of software written for the central unit. Adding a new feature does not require building a single extra wire or connection, just to add a few lines to the software.

Center of Attention


The four networks all connect to the central unit, the **Built-in Systems Interface (BSI)**. This control unit manages the flow of information between the networks (many of the

messages generated in one network has to be relayed to another, just one example is the suspension computer—connected to CAN—being interested in messages about the open or closed position of the doors—communicated on VAN 2).

In addition to that, the BSI offers an interface to the outside world as well, a diagnostic socket which can be used to check, test and configure the whole system.

The multiplex system switches to an energy-saving low power mode whenever possible.

Air Conditioning



Air conditioning

Once considered pure luxury, air conditioning and other forms of climate control have become standard items on today's car. After all, creating an acceptable environment for the driver is more than a mere question of comfort, it contributes to safety to a great extent.

XM

There were several **climate control** systems fitted to our Citroëns, offering various degrees of automation of keeping the climatic conditions inside the car. The system can be **manual**, **semi-automatic** or **automatic**. The manual version also came with separate settings for driver and passenger.

The semi-automatic system is rather similar to the manual one, the visible difference is that the operating knob on the dashboard is marked in degrees instead of just blue and red. The direction and recirculation controls are identical to the manual system. The automatic climate control looks radically different, with a controlling panel using buttons and a digital temperature display.


The AC system in the XM is fairly simple. If it is on, the air is always cooled to about 8–10 °C on the inlet side (this is varied between the air intake from outside and recirculation from the inside) and then if you set a higher temperature, it's reheated. The heater also always works, its effect is only regulated by allowing air to flow or not to flow through it (this is what the flap valve does). The air always flows through the AC heat exchanger. As a result, the AC also dries out the air whenever it is on. Once the air passes out of the temperature regulating flap valve, another flap valve regulates where it goes inside the cabin. That's really all there is to it.

The AC system itself is almost self-sufficient. It has a radiator, compressor, heat exchanger with evaporator, and a condenser—and the connecting pipes. The climate control ECU actually only provides a signal to a relay that switches the AC system on by operating the electric clutch on the compressor. This same signal switches the radiator fans on to the low speed. The AC system in turn sends a 'fans to full speed' signal to the fan controller, when the coolant temperature reaches a trip point (this is handled by a different switch section in the same pressure sensitive switch that prevents the AC going on without any coolant in the system, described above).

As far as I know (unless it changed in later versions), the AC itself (as opposed to climate control) never had an ECU.

The **evaporator** has an integrated pressure/temperature valve, opening up the pressure line to the return line.

After coming through the evaporator, the temperature of the fluid (more precisely, a mixture of liquid and vapor) suddenly drops because of the drop in pressure. It enters the **heat exchanger** which operates like a radiator, cooling the air and heating itself up. The fluid then goes back to the drier-radiator-compressor end of the loop. The con-



densed moisture is collected from the heat exchanger and let out through to floor of the cabin via a plastic tube.

As the air always enters through the heat exchanger, and whether it gets cooled at this point, depends only on whether the compressor is working or not. The temperature flap only decides which part of the air is going to be taken before or after the heater radiator. This is how the temperature is regulated.

When the **compressor** is on, is to condense the moisture out of the air, and then re-heat it as necessary to the temperature set on the controls. Since the temperature is regulated by the temperature flap, it has really nothing to do with the compressor at all—the only consequence of the compressor not working (for any reason) is that the system will obviously not be able to produce a temperature lower than ambient.

There are four sensors providing input. The first one is at the entrance of the air, before the heat exchanger, the second one after the temperature flap, the third one on the roof, and the last one in the heat exchanger. They have very different but sometimes overlapping roles.

The first three collectively influence temperature regulation. In particular, the sensors after the temperature flap and on the roof determine what the actual temperature is. The sensors before the heat exchanger and after the temperature flap decide how fast the temperature flap will be moved to prevent extremely fast changes in temperature in the cabin. This does not always work very well, which is why you get a blast of air when the system is set on auto and you leave the car in the sun in summer. Both of these parameters (temperature and temperature difference) influence the fan speed.

The sensor before and in the heat exchanger as well as the temperature selection, influence the AC part, i.e. the operation of the compressor. For instance, the compressor will not operate below a certain external temperature. Also, it will not operate if the temperature is set to maximum.

When the system is cooling the incoming air, it needs to have the exchanger at a temperature which is lower than the ambient air temperature, obviously. As the compressor either runs or not, it cannot cool just a little bit—it always either runs on full or does not run. When it starts, it will start cooling the heat exchanger. How cold it will get, depends on how hot the incoming air is and how much air is coming in. In any case, when it gets significantly colder than the incoming air, the moisture from the air will start to condense on the heat exchanger, which is why there is a collector underneath it and a drip outlet. If the compressor keeps on working, while the heat that needs to be taken from the air is lower from the heat transfer ability of the whole system, the heat exchanger will continue to progressively get colder. If nothing is done, it will get well below freezing (it can go as low as -40°C given proper fluid, and of course construction designed for this). What will happen then is that the condensed water from the air will start freezing on the heat exchanger fins, and eventually, the whole thing will become a solid block of ice (usually there will be a crackling noise to accompany the event), preventing actual air flow. If the condition persists, the pressure in the system will build up until the valve in the evaporator opens, and by this time it is possible that the fluid actually gets heated up enough that the remaining part going through the heat exchanger will actually melt the ice producing a fog (I've seen it happen!). All of this will be the lucky turn of events, assuming the ice has not cracked the heat exchanger and that there is no fluid leak.

So, obviously, there is a sensor, and that's the fourth one in this story, which detects the temperature of the heat exchanger becoming too low. When that happens, the compressor is cut out, until the heat exchanger temperature rises to an acceptable level. The thermal inertia and different cut out and cut in temperatures insure that the compressor doesn't keep switching on and off too quickly, which would place an undue strain on the electromagnetic clutch.

The logic in the ECU is done very simply, if the fourth sensor detects that the heat exchanger is too cold, the compressor will switch off, regardless of the AC switch and temperature set. The only thing it will do, as I said in the earlier mail, is that it will switch on for about 1 second whenever the AC switch is turned on, this is probably some ECU feature. The compressor will never turn on if the gas pressure is insufficient, and this part is handled by the pressure switch on the drier, and has nothing to do with the ECU. In fact, the ECU only gives the whole system a 'go-ahead'.

Appendix



ORGA number

This number shows the day when your car was actually assembled on the production line. The dealers and parts stores use this number (often called ORGA or RP number, the second stands for Replacement Parts) to identify the various parts and components fitted to your car.

On various models, the ORGA number can be found in different locations. It is on the top of the left hand suspension turret on **Visas**, **C15s**, **AXs** and **CXs** (often hidden by the wiring harness). **BXs** and **XM**s have it stamped on the left hand front door A-pillar, above the courtesy light switch. On the **GSA** you find it on the inner right front wing. **Xantia**e switched to the other side: the number can be found on the bulkhead just in front of the right suspension sphere. C5 ????????????

Calculating the production date is very easy using the following table. Locate the largest number in the table still less than or equal to your organization number. To see an example, let's assume the number is **4859**. Then the largest number will be **4832** in the cell **February 1990**. Just subtract this number from your organization number to get the day of the month of the production of your car (in our example, $4859 - 4832 = 27$ yields **February 27, 1990**).

If you receive the non-existent date zero (this happens when your organization number is not greater than but equal to the number in the table), simply take the last day of the previous month. For instance, for the organization number **5013** the largest number in the table is **5013** in the cell **August 1990**, subtraction results in **zero**, hence the production date is **July 31, 1990**.

Years	Months											
	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
1982	1879	1910	1938	1969	1999	2030	2060	2091	2122	2152	2183	2213
1983	2244	2275	2303	2334	2364	2395	2425	2456	2487	2517	2548	2578
1984	2609	2640	2669	2700	2730	2761	2791	2822	2853	2883	2914	2944
1985	2975	3006	3034	3065	3095	3126	3156	3187	3218	3248	3279	3309
1986	3340	3371	3399	3430	3460	3491	3521	3552	3583	3613	3644	3674
1987	3705	3736	3764	3795	3825	3856	3886	3917	3948	3978	4009	4039
1988	4070	4101	4130	4161	4191	4222	4252	4283	4314	4344	4375	4405
1989	4436	4467	4495	4526	4556	4587	4617	4648	4679	4709	4740	4770
1990	4801	4832	4860	4891	4921	4952	4982	5013	5044	5074	5105	5135
1991	5166	5197	5225	5256	5286	5317	5347	5378	5409	5439	5470	5500
1992	5531	5562	5591	5622	5652	5683	5713	5744	5775	5805	5836	5866
1993	5897	5928	5956	5987	6017	6048	6078	6109	6140	6170	6201	6231
1994	6262	6293	6321	6352	6382	6413	6443	6474	6505	6535	6566	6596
1995	6627	6658	6686	6717	6747	6778	6808	6839	6870	6900	6931	6961
1996	6992	7023	7052	7083	7113	7144	7174	7205	7236	7266	7297	7327
1997	7358	7389	7417	7448	7478	7509	7539	7570	7601	7631	7662	7692
1998	7723	7754	7782	7813	7843	7874	7904	7935	7966	7996	8027	8057
1999	8088	8119	8147	8178	8208	8239	8269	8300	8331	8361	8392	8422
2000	8453	8484	8512	8543	8573	8604	8634	8665	8696	8726	8757	8787
2001	8819	8850	8878	8909	8939	8970	9000	9031	9062	9092	9123	9153
2002	9184	9215	9243	9274	9304	9335	9365	9396	9427	9457	9488	9518

Index



Index

AAV	8	DIRASS	25, 42
ABS.	25, 50	DIRAVI	44-45
computer.	50	direct injection	19
hydraulic block	50	distributor	
sensor	50	diesel.	12
AC	56	ignition	9
accelerator pedal	13, 17	door/tailgate open sensor	29-30, 34
Activa	36-37	EDC	17
computer.	37	EFL	6
electro-valve	36	EGR.	16, 19
roll corrector	36	electronic diesel control	17
sphere	36	EMS	9
stabilizing cylinders	36	engine management system	9
AFS	6, 8	engine runaway	14
air conditioning	56-57	evaporator	56
air flow sensor	6	exhaust gas recycling	16
air temperature sensor	7	flow distributor	42
anti-dive behavior	23	fuel cut-off	
anti-roll bar.	24	diesel.	18
anti-sink	35	gasoline	8
sphere	35	fuel filter	
valve	35	gasoline	8
ATS	7	fuel injection	
autotransformer	9	diesel, direct	19
auxiliary air valve	8	diesel, electronic	17
ball and piston valve.	27, 32	diesel, mechanical	12
BHI	38	gasoline, electronic	6
body movement sensor	29-30, 33	fuel pump	
brake		diesel.	12
compensator valve.	26, 48-49	gasoline	8
force limiter	49	fuel stop valve	12
front cylinders	26	fuel supply	
pressure sensor	29-30, 34	gasoline	8
broadcast messages	53	fuel tank	
BSI	53	gasoline	8
Built-in Hydroelectronic Interface	38	glow plug	15
Built-in Systems Interface	53	ground clearance	23
bus	52	HDI	19
CAN	38, 52	heat exchanger	56
CAS	9	heated wire AFS	6, 17
catalytic converter		height corrector.	26, 29, 32
diesel.	18	front	26
gasoline	10	high pressure pump	
chip tuning	6	diesel.	12
climate control	56	hydraulics.	25, 38, 43
cold start injector.	8	HP brakes	48-49
compressor.	56	Hydractive 3.	38-39
Controller Area Network	52	Hydractive I	27, 29, 31
coolant temperature sensor	7	Hydractive II.	32-33
crank angle sensor	9	Hydractive sphere	27, 32
CSV	8	Hydractive valve	
CTS	7-8	electric	27, 32
damping elements.	26-27, 32	hydraulic	27, 32
DI	19	ICSM.	7
diesel combustion	12	idle control stepper motor	7

Idle speed	7	wheel centering device	45
idle speed control valve	7	wheel speed sensor	29-30, 33
ignition coil	9	stiffness regulator	39
ignition delay		suspension ECU	27-29, 32, 38
diesel	14	suspension mode	
gasoline	9	hard	28-29, 32, 39
ignition key switch	7	soft	28-29, 31-32, 34, 39
ignition switch	29, 31, 34	suspension resonance frequency	23
injection adjuster	12, 14, 18	suspension selector switch	28, 31-32
injection delay		suspension status light	31
diesel	14	swirl chamber	12
gasoline	8	TDC	9
injector		temperature-timer switch	8
diesel	15	throttle pedal position sensor	29-30, 33
gasoline	8	throttle position switch	7
injector needle movement sensor	18	throttle potentiometer	7
inlet manifold	8-9	timing advance	9
intercooler	15	top dead center	9
ISCV	7	TP	7
knock sensor	10	TS	7
KS	10	turbocharger	
lambda		diesel	15
ratio	6, 10, 18	VAN	52
sensor	10	variable turbo pressure	18
LDS	38	Vehicle Area Network	52
lead in gasoline	10	vehicle speed sensor	29, 33
leak returns	25	wastegate	
main accumulator	25, 42	diesel	15
manifold absolute pressure	6		
MAP sensor	6		
monopoint EFI/EMS	8		
multiplex message	53		
multiplex wiring	52		
multipoint EFI/EMS	8		
ORGA number	60		
OS	10		
oxygen sensor	10		
particulates	14, 18		
PAS	25, 42		
piston and ball valve	28		
post-glowing	16		
power assisted steering	25, 42-43		
power steering	25, 44		
prechamber	12		
pressure regulator			
EFI	8		
hydraulics	25, 42		
regulator	12-13, 17		
reservoir			
hydraulics	25		
RP number	60		
security valve	26		
self-diagnostics	18		
smoke limit	14, 18		
spark plug	9		
steering			
centering pressure regulator	45		
control unit	45		
control valve	42		
rack	45		
ram cylinder	42, 45		
wheel angle sensor	29-30, 33		