PRINCIPLES of PNEUMATICS

TRADE & TECHNICAL PRESS LTD.

PRINCIPLES AND THEORY of PNEUMATICS

England
Trade & Technical Press Ltd.
Crown House, Morden, Surrey

PREFACE

EVERY ENGINEER connected with the design or utilisation of pneumatic equipment should possess in some handy form the essential information which will enable him to deal promptly with everyday problems likely to arise relating to pneumatic engineering, and this volume aims to achieve that purpose. "Principles and Theory of Pneumatics" is presented in a manner which permits a quick reference to the basic principles of the industrial users of pneumatics. Engineering students will also find this book equally valuable in their studies.

THE PUBLISHERS

CONTENTS

SECTION 1

Sub-Section A

Principles and Theory of Pneumatics

INTRODUCTION

PNEUMATICS—derived from the Greek "pneumos" meaning breath—and always treated as a singular noun, deals with the behaviour of gases, and we are mainly concerned with one branch of this, namely, the use of gases for transmitting power.

All gases are readily compressible and it is this property which differentiates them most from liquids as a power transmission medium. The behaviour of a gas when transmitting power can be seen very easily with an ordinary bicycle pump. If the handle is pulled out and the outlet covered with the finger the air inside it will behave very like a spring, and a weight placed on it will bounce up and down. If a fairly heavy weight is put on a table and pushed along by the extended pump still with the outlet stopped, it will be noticed that the pump goes in and out as the friction of the weight on the table varies.

Bouncing the pump handle up and down will not produce any significant heating, but if the pump is used continuously to force out air under pressure, it will get quite hot, as will the air leaving it. On the other hand, if a bicycle tyre is let down, the air leaving it will feel quite cold—colder even than the draught it causes. It may even make the valve so cold that it acquires a coating of ice.

The compressibility of a gas was first investigated by Robert Boyle in 1662 and he found that the produce of pressure and volume of a particular quantity of gas was constant, providing the temperature did not vary.

This is usually written as—

$$PV = C \text{ or } P_1V_1 = P_2V_2$$

In this equation the pressure is the absolute pressure, which for "free" air is about 14·7 psi and is, of course, capable of maintaining a column of mercury nearly 30 inches high in an ordinary barometer. Halving the volume of a cylinder by pushing in the piston half-way, increases the absolute pressure sufficiently to sustain a mercury column nearly 60 inches high.

An ordinary pressure gauge shows the pressure above atmosphere and the readings are given as psig, i.e. pounds per square inch gauge. To convert "gauge" pressures to absolute pressures, it is necessary to add 14·7 psi. An ordinary shop compressor would deliver at 100 psig or 114·7 psia. For the purposes of calculation it may sometimes be convenient to express pressures in atmospheres, when 100 psig equals 7 atm(g) or 8 atm(a).

Any gas can be used in a pneumatic system but air is the most usual, for obvious reasons. Exceptions are most likely to occur on aircraft and space vehicles where an inert gas such as nitrogen is preferred, or the gas is one which is generated on board.

Air is a mixture of 78% nitrogen, 21% oxygen by volume; water being the most important remaining ingredient as far as pneumatics is concerned. The dilution of the oxygen by nitrogen makes air much less chemically active than pure oxygen, but it is still capable of causing spontaneous combustion or explosion, particularly if oil vapour at an elevated temperature is present, as may occur in an air receiver.

One cubic foot of air at standard temperature and barometric pressure weighs 0·0765 lb. and can contain ·0011 lb., or about 2%, of water at 100% humidity. Its specific heat is 0·242 or about one quarter that of an equal weight of water.

The Air Power Transmission Cycle

Energy is expended in compressing a gas a proportion of which appears as heat and usually performs no useful purpose—it may require costly plant to remove it—whilst the remainder is available as potential energy. The compressed air is conducted through pipes to the point of application, where it performs useful work, either by expansion, or by direct force. It is then exhausted to atmosphere.

Compressed air can be stored for use when demand arises or is produced as required. The former can meet only small demand and is more likely to be found in conjunction with a hydraulic system, where the air is compressed to several thousand pounds per square inch. In the ordinary distribution system the storage capacity given by the air reservoir and piping gives a certain amount of flexibility, so that it is not necessary for demand and supply to match exactly at any particular instant as in a hydraulic system, although they must do over a short period. As will be shown later, the normal storage capacity can only deal with short term fluctuations of demand.

To many users the advantage of compressed air as a source of power is that it can be plugged into and used to produce reciprocating and some kinds of rotary motion at comparatively little cost and with very little trouble.

Normal Pressures

A nominal working pressure of 100 psig has been accepted for most workshop air systems, which in practice means that the pressure at the point of application may vary between 80 and 100 psi. Any sudden demand, outside the compressor capacity, is met by lowering the pressure in the system either as a whole, or if the pipes are not large enough, in the locality of the demand.

This pressure allows compressors to be comparatively simple whilst allowing the cylinders etc. to be kept to reasonable dimensions.

Although compressed air is potentially dangerous, accidents are unlikely at 100 psi if the normal precautions are taken. Higher pressures are only justified where the expense of the special equipment justifies the savings in weight and size, as on aircraft, where pneumatics may compete with hydraulics on a power/weight basis.

Air Compression

Air compressors are highly developed machines but their design is by no means finalized. For many years piston compressors were in the vast majority, but their predominant position is being seriously threatened by vane and screw types in the 100 psi field. All have the same basic principle and the differences are chiefly because of mechanical and heat transfer considerations.

A gas can be compressed isothermally (Fig. 1) when the heat of compression is extracted as quickly as it is produced, or adiabatically (Fig. 2) when the heat remains constant or, in practice, by an intermediate condition when only a proportion of the heat is extracted.

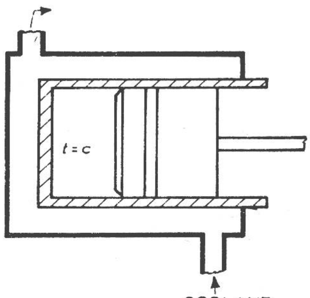

FIG. 1.
Compressing air isothermically, where the temperature is kept constant by heat passing to the coolant as it is generated.

To compress a given volume of air isothermally the horsepower is given by:—

$$\text{H.P.} = \frac{144 \times P_1 \times \text{F.A.D.}}{33000} \log_e \frac{P_2}{P_1} \qquad (1)$$

where P_1 is initial pressure psi (abs.)
P_2 is final pressure psi (abs.)
F.A.D. is free air delivered cu. ft./min.

For adiabatic compression:—

$$\text{H.P.} = \frac{144 \times P_1 \times \text{F.A.D.}}{33000} \frac{\gamma}{\gamma - 1} \left[\left(\frac{P_2}{P_1} \right)^{\frac{\gamma-1}{\gamma}} - 1 \right] \qquad (2)$$

where γ = ratio of specific heats at constant pressure and volume $= 1 \cdot 4$.

The power required for adiabatic compression at 100 psi (gauge) is about $1 \cdot 4$ times that for isothermal compression, so that the nearer the conditions are to isothermal the more efficient the compression.

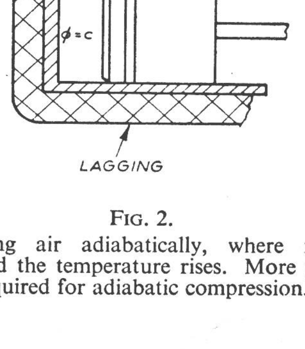

FIG. 2.
Compressing air adiabatically, where no heat escapes and the temperature rises. More power is required for adiabatic compression.

In a piston compressor heat is extracted through the cylinder walls. With a modern vane compressor most of the heat goes to the lubricating oil which is injected into the compression zone and passed to a cooler after separating from the air.

The theoretical power to compress 100 cu. ft. of free air per minute to 100 psig as calculated from (1) and (2) respectively, are:—

Isothermal 13·2 H.P.
Adiabatic, single stage 18·0 H.P.
Adiabatic, two stage 15·4 H.P.
Actually measured for piston compressor 22·1 H.P.

The quantity of air delivered by a compressor or consumed by a tool is always given in cu. ft. of free air, i.e. air at 29·9 in. barometer and 20°C. This is often known as F.A.D. The volume of air, when compressed, is the free air volume divided by the absolute delivery pressure in atmospheres. One cu. ft. of free air compressed and delivered at 100 psig requires the expenditure of about ¼ h.p. at the compressor, being more for small compressors and less for large ones.

Multi-stage Compression

It is obvious that a lot of power can be saved by approximating to isothermal compression, and this is done in practice by multi-stage compression.

The air from the first, low pressure, stage is passed through an inter-cooler which may consist of a finned pipe over which cold air is blown by a fan on the flywheel on a small compressor, or a water cooled heat exchanger on a larger one. It reaches the second stage cylinder comparatively cool and is then again heated up, so that a further intercooler is used if there are more than two stages.

The arrangement of a two-stage compressor is shown in Fig. 3 and this would give a power saving of about 15% as compared with a similar single stage compressor.

The use of multi-stage compression depends on the size of compressor. On the smallest, pressures up to 400 psig are obtained in a single stage, as there is no economic justification for complicating the compressor for the small power saving which would be obtained.

On larger installations, running costs become of prime importance and warrant elaborate water cooling both of the compressor cylinders by water jacketing and of the intercoolers. It is usually considered that two-stage compression is desirable for pressures over 90 psig and three stage for pressures over 300 psi. For very high pressures more stages are necessary.

Piston Compressors

Piston compressors are of many types—single or multiple, single or double-acting, high speed, low speed, air or water cooled, but the variations are matters of convenience or to suit the ideas of the designer.

For the most efficient compression in any type, the cylinder must be kept as cool as possible, the dead space at the top of the stroke kept to a minimum and the resistance to air flow into the cylinder negligible.

Valves which lift easily and require little depth have been specially designed for compressors. Some makers, however, fit a cam-operated inlet valve which increases the inlet flow efficiency.

Compression takes place with a partial loss through the cylinder walls of the heat generated, i.e. it is neither isothermal (at the same temperature) nor adiabatic (without transfer of heat). Under these conditions the pressure volume relationship is given by:—

$$PV^{1\cdot3} \text{ (approx.)} = \text{constant}$$

Piston compressors are controlled either by stopping and starting, by lifting the inlet valve mechanically or by throttling the inlet so that the cylinder is only partially filled.

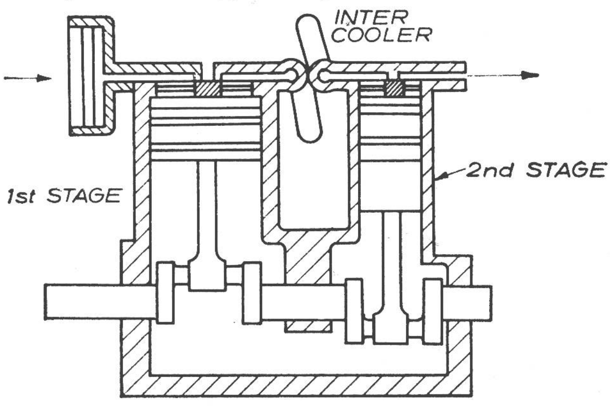

FIG. 3.
Two stage compression. The larger low pressure piston draws in air and forces it through the intercooler to the high pressure cylinder. The intercooler can be air or water cooled and the cylinders themselves are often cooled too. The inlet and exhaust plate valves are fitted to the cylinder heads.

When stopping and starting under pressure the compressor must be "unloaded" by opening the inlet valve until the compressor is running at full speed when operation proceeds normally.

Vane or Rotary Compressors

The principle of the vane compressor is essentially simple as it is valveless and has no pistons. The practical difficulties centre on the sealing of the vanes and this is now usually done by flooding the compressing cylinder with oil which also dissipates a large proportion of the heat of compression.

The cycle of compression is shown in Fig. 4. Air is entrained between the vanes for about 90° and is then compressed as the space between the rotor and casing diminishes. After compression is about 60% complete, oil is injected into the space, which absorbs a lot of the heat generated by the air and also reduces the volume available for it, so that the final compression is very close to isothermal.

Some of the practical features of a vane compressor are shown in Fig. 5.

The outer casing is pressurized at the discharge pressure and oil can therefore be delivered through a pipe from the oil sump to the cylinder at a zone where compression is only partly complete, so making an oil pump unnecessary. The mixture of air and oil is discharged into the outer casing and the air then passes over a series of baffles which separate out most of the oil, so that it drops back to the sump. Any remaining oil is removed by the fabric filters in the discharge stream and drawn back through a pipe to the suction zone. It is claimed that the discharged air may contain less oil than that from a piston compressor.

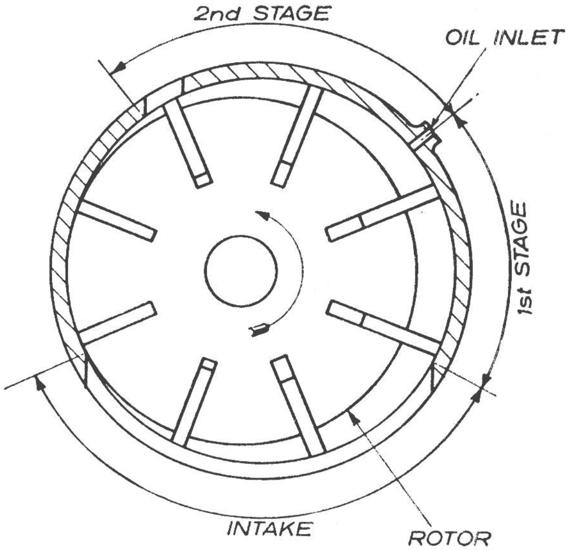

FIG. 4.

Diagram of vane compressor with oil injection. The air is trapped between the rotor and casing by the vanes and compressed as the space narrows. Because of the injection of cold oil, the compression is essentially equivalent to a two-stage piston compressor.

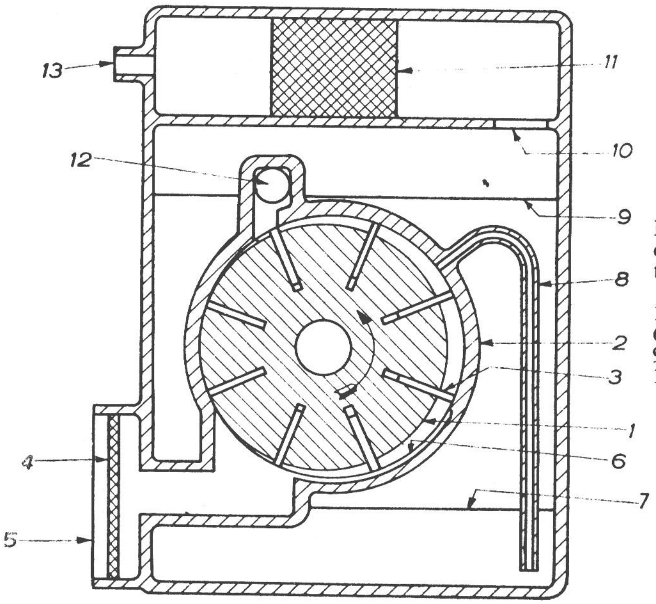

FIG. 5.

Diagram of vane compressor showing some of the essential features. Features not shown are inlet throttle to control volume of air delivered, oil return pipe from oil filter to intake and oil cooler.
1. Rotor 2. Drum 3. Vane 4. Inlet filter 5. Inlet 6. Inlet groove 7. Oil level 8. Oil injection pipe 9. Baffle for oil separation 10. Air passage to filter 11. Oil filter 12. Outlet to pressure chamber 13. Outlet

The oil, which is responsible for practically all heat dissipation, may be cooled by passing it through finned pipes on to which air is blown by a fan mounted on the motor shaft.

Large vane compressors may have water cooling.

The output of a vane compressor is practically pulseless and this reduces the need for an air receiver. By fitting a pressure sensitive servo regulator operated by oil pressure from the lubricating system it is possible to control the output to balance the input.

AIR COMPRESSOR CONTROL

It is normal practice for sufficient air compressor capacity to be installed to give an appreciable margin over the average demand, so as to allow for extension of the system and for peak demands. Because of this it is usually necessary to provide a control method to balance the compressor output with consumption. It is not practicable and certainly very uneconomic to blow off excess air through a relief valve as in a hydraulic system.

Compressor output is, in practice, controlled by the pressure changes at the compressor outlet or in the receiver, if one is fitted. As far as is known, flow measurements are not used. The accuracy of control can be as good as the pressure sensing control device makes possible, although in ordinary workshop installations quite large pressure variations are often tolerated.

Compressor output can be regulated in the following ways:—

1. Stopping and starting
2. Unloading—by lifting the suction valves
3. Varying speed
4. Regulating the admission of air at the inlet

Stopping and Starting

This method has much to recommend it, especially where several compressors are needed to carry the peak load. A large compressor may require 20% of full load power when running idle, and on a small compressor, it may use to well over 50%. Most of this waste power has to be dissipated in the cooling system.

The number of starts per hour should not exceed 15 or 20 an hour or there may be excessive electrical maintenance costs. The frequency of starts with a single compressor varies with the relation between compressor size and demand. The greatest number is when the compressor capacity is double the demand. If it is less, it is running for longer than it is stopped and if it is more it is stopped longer than it is running. The most influential factor, however, is the size of reservoir in relation to demand. The larger the reservoir and other system capacity, the smaller the number of starts.

To take as an example a small compressor with an output 100 F.A.D. and a demand of 80 F.A.D. The receiver and pipe line capacity is 50 cu. ft. and the allowable pressure variation 15 psi.

When the compressor starts at the lower pressure it will meet the demand of 80 F.A.D. and have an excess of 20 F.A.D. which passes to the receiver, which will recharge it in $\frac{50}{20}$ mins. or $2\frac{1}{2}$ mins. This will meet the demand for rather less than a minute, so that the compressor will be working on a 3 min. cycle, or 20 starts per hour.

If the demand increases to 90 F.A.D. there will only be 10 F.A.D. for the receiver, which will take 5 mins. to charge and then discharge in about $\frac{1}{2}$ min. giving a $5\frac{1}{2}$ min. cycle. Reduction of the demand to 60 F.A.D. gives a receiver charge time of $\frac{50}{40}$ or $1\frac{1}{4}$ mins. and a discharge time of $\frac{50}{60}$ mins. or 5/6 mins. or 2 1/12 min. cycle.

As a 50 cu. ft. air reservoir measures 3 ft. dia. $\times$ 7 ft. long or 2 ft. 6 in. dia. $\times$ 10 ft. long, there will probably be some reluctance in purchasing anything larger for a 100 F.A.D. compressor, taking only 25 H.P.

One method that has been used for keeping down frequency of starts without increasing the receiver size, or with too wide a pressure fluctuation, is to fit a reducing valve between the receiver and distribution pipes. By this means the air enters the system at an almost constant pressure, whilst the reservoir pressure may fluctuate 20 or 30 psi. There is, however, a noticeable loss of energy.

Compressors controlled by the "start and stop" method are almost invariably electrically driven. For a small compressor a simple pressure switch (Fig. 6) is piped to the air receiver and this may operate the motor direct when direct-on-line starting is permissible. For larger powers it would control a contactor starter.

Controlling a number of Compressors

The stopping and starting method of control is most useful for large installations which are supplied from several compressors. The number of compressors running at any one time is only

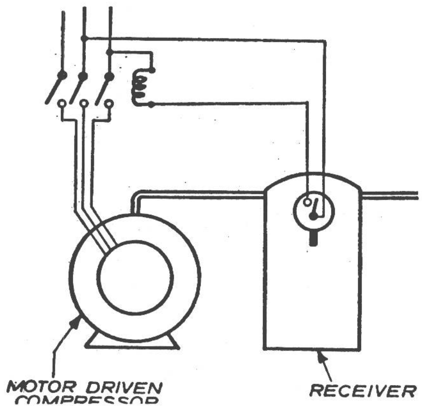

FIG. 6.
On/off control of compressor by pressure switch. The switch inlet must be connected to the receiver to protect it from pressure pulses from the compressor.

sufficient to meet the load, and small fluctuations are met by stopping and starting one or more of the compressors as required. Normally one or more compressors are selected to take the base load and run continuously whilst the demand exists.

To do this each compressor is controlled by its own pressure switch, the switch being set at different pressures. For example, if there were three compressors (Fig. 7) one might be set at 105 psi "out" and 100 psi "in", the next 100 psi "out" and 95 psi "in", the third 95 psi "out" and 90 psi "in". All three compressors would be started up initially and as the pressure rose the first would cut out at 95 psi and if the two remaining machines were more than equal to the demand, the second would cut out at 100 psi. Should the demand drop sufficiently, as at the end of the shift, the pressure might rise to 105 psi as the last compressor stops. As the pressure falls, through leakage or demand, the compressor which cut out last will be the first to start. It will be recalled that these figures would only be approximate with commercial pressure switches and would need adjustment to suit the local requirements.

Unloading

To make starting easier, piston compressors are fitted with an unloading device which prevents the suction valve seating and so stops it operating until running at full speed. Use can be made of this device for matching the compressor output with the demand and it is, in fact, almost a necessity with engine driven piston compressors where the engine cannot be stopped.

The unloader is pressure operated and is shown in Fig. 8. A spring-loaded valve—similar in principle to a safety valve—is connected to the receiver and open when the pressure reaches the figure to which the valve has been set. Air then flows to the piston of the unloader, which lifts the inlet valve off its seat. The air in the cylinder, instead of being compressed, blows back through the valve. When the pressure falls, the unloader valve shuts and the air in the unloader leaks away.

It is most important that no dirt enters the valve and an effective filter is essential. This can be made of felt or porous metal.

On electrically driven compressors, the unloader is controlled by a solenoid operated three-way valve energised from the motor starter which admits air from the receiver to a plunger which engages the suction valve (Fig. 9).

Control by Speed Variation

This is not a common method of compressor control because of the complications involved. The ideal is to match the compressor output with the demand within narrow pressure limits by varying the speed of the compressor shaft.

An electrical drive involves a variable speed motor which in practice would almost certainly mean a Ward Leonard Drive, involving a DC dynamo driven by an AC motor and a variable speed DC motor.

An I.C. engine drive would be easier to

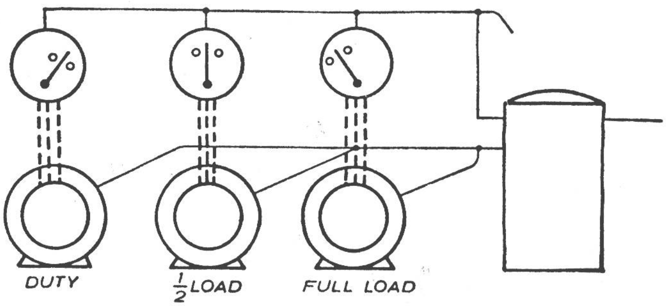

FIG. 7.
On/off control of three compressors arranged to select the number running to suit the load.

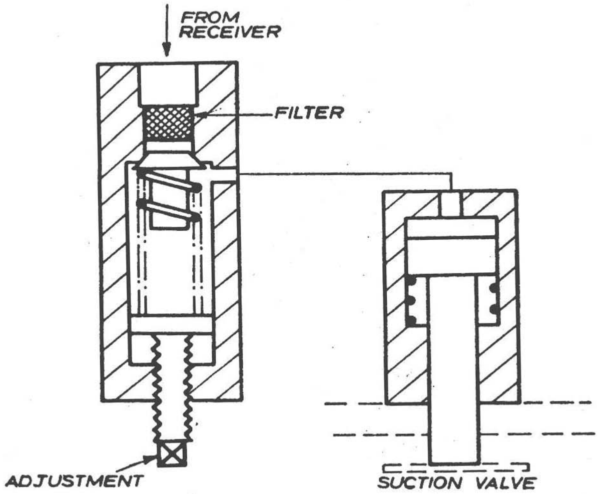

FIG. 8.
Pressure operated unloader. When the pressure becomes too high, it overcomes the spring, opens the mitre valve and so applies pressure to the plunger which depresses the suction valve on the compressor.

control, but as the effective speed range is limited there are obvious drawbacks.

As far as is known, potentially the most satisfactory method—a fluid power drive—has not been used. This would be arranged as shown in Fig. 10.

A constant speed AC electric motor drives a hydrostatic transmission consisting of a variable delivery pump and fixed capacity (i.e. constant torque) hydraulic motor which in turn is coupled to the compressor shaft. Theoretically it should now be possible to control the compressor speed by keeping the pressure in the hydrostatic transmission constant. Any rise in air pressure due to falling off in demand, tends to increase the torque, and the reaction on the hydraulic pump control would reduce its delivery rate, so slowing down the motor.

If the pressure variations, inherent in this rather remote method of control are unacceptable, it would be possible to use a pneumatic pressure controller which would operate the hydrostatic speed drive control through an air cylinder. A modern pressure controller would give very exact pressure control.

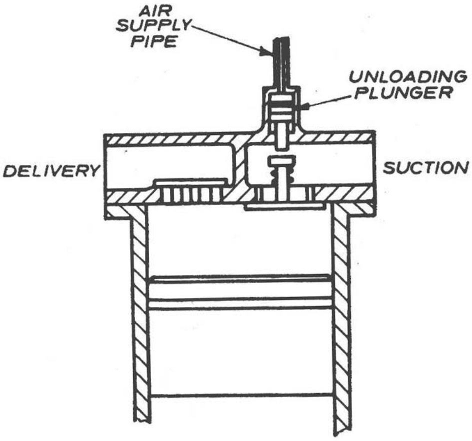

FIG. 9.
Unloader for electrically driven compressor. When the motor starts a special switch on the starter energises a three-way valve to pressurise the unloader plunger. The switch cuts out when the motor is running at full speed.

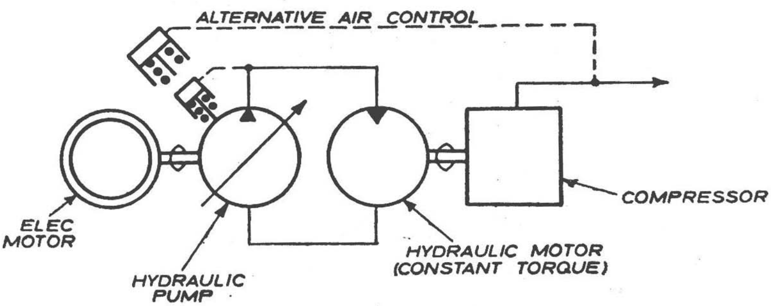

FIG. 10.
Variable speed hydraulic compressor drive which maintains compressor output pressure constant, irrespective of load.

Regulating at the Suction Inlet

A method used very successfully on modern vane compressors involves restricting the amount of air drawn into the machine in accordance with the demand. With this method the pressure acts on a spring-loaded pilot valve which controls the air supply to an unloading valve in the inlet. This type of control is very effective and maintains the pressure constant to about 1 psi as long as the demand is within the capacity of the machine.

Overloading

It is not possible to overload an air compressor by excessive demand, although it would be possible to do so if all outlets were shut and the control mechanism did not function. A safety valve capable of dealing with the full compressor capacity is necessary to deal with this eventuality.

If the potential demand is greater than the compressor can supply, the first symptom will be a fall in pressure and a corresponding reduction in the power taken by the compressor, so that the power available actually becomes less rather than more. This condition is therefore readily recognised by the pressure gauge and ammeter readings being below normal.

Flow of air in Pipes etc.

In practice, the pipes of a compressed air distribution system fulfil two functions—to convey the air to the point of use and to act as reservoirs to meet local demands. Although pipe friction causes loss of power, its most objectionable feature is that it causes local demands to produce a pressure drop which may affect other users in the vicinity.

Flow of compressed air, as with any other confined fluid, can only take place when there is a drop of pressure along the pipe. When flow does take place there is a friction loss proportional to the square of the velocity.

If we have a pipe leading from a receiver to a distant point, there will be no flow in the pipe whilst all valves are shut, but the pressure throughout the length of the pipe will be the same as that in the reservoir—say 100 psig.

If a valve is opened feeding a tool taking 20 cu. ft. free air /min. there will be an immediate pressure drop which will start the flow of air from the receiver. By reference to the tables for pressure drops the figure to which the pressure will fall can be obtained, given the size and length of supply pipe. For example, 500 ft. of ½ in. pipe would cause a pressure drop of 9 psi. This is not serious, but if a second similar tool is connected, then the demand is doubled, but the friction loss, varying with the square of the velocity, jumps not to 18 psi but to 36 psi, which is obviously far too great for the tools to work satisfactorily.

Because this is not recognised, air lines at the extreme of a system are often overloaded. It cannot be too highly stressed that double the consumption causes the friction to be quadrupled and by tripling it, it is increased ninefold.

One method of mitigating the effects of pipe friction is the ring main, by which any demand is fed from two directions.

Although pipes can act as reservoirs to meet momentary demands from cylinders, when the demand is continuous as with a cutting tool, then the pipework must be capable of supplying the full demand without an unacceptable pressure drop.

Air Cooling

Whilst inter-cooling between the stages of compression is intended to improve efficiency, the principal reason for after cooling is to remove as much moisture as possible before the air passes to the distribution network. Free water in an air system is most undesirable as it can cause water hammer, rusts and chokes valves and is generally objectionable.

The amount of moisture which air can absorb depends only on the volume and the temperature and is unrelated to pressure. When air at 50% humidity is compressed therefore to, say, one eighth of its original volume, it becomes saturated (100% humidity) at 2 atmospheres, assuming the temperature remains constant, and thereafter the remainder of the water content is deposited as liquid, so that finally three quarters of the original moisture content has been deposited. In practice, of course, the air is heated as it is compressed and is so able to absorb more water than when it cools. If it is cooled as it leaves the compressor, most of the moisture can be removed before it passes to the pipes; if it is still warm the moisture will form as dew on the pipe walls and will get carried along with the air stream.

For small compressors, natural cooling is usually considered sufficient but even here some attempt at cooling, such as running the pipe along the outside of a north wall, will make an appreciable difference to the amount of water which can be collected.

Larger installations usually warrant an after-cooler, which is a tubular heat exchanger with water passing upwards through the tubes and the air flowing downwards. At the bottom a float trap discharges the deposited water and oil. As oil acts as an efficient insulator the air side must be kept clean if the cooler is to maintain its efficiency. Mains water is used where possible and a temperature rise of the water of 20°F. is usually accepted as economical.

Air pre-heating

The classic example of pre-heating compressed air prior to use is in the Whitehead torpedo where the compressed air, before passing to the engine, is heated to about 300°F.

The main reason for pre-heating on land installations is to prevent excessive low temperature due to cooling when expanding, with the formation of ice. On the torpedo application, however, the heater is used to augment the energy stored in the compressed air.

If waste heat is available, then it can profitably be used for heating air, providing the heater is immediately prior to the point of use. With waste steam a tubular heat exchanger is the most suitable form of heater.

Air Distribution and Moisture Removal

After leaving the compressor and aftercooler, if any, it is advisable to remove the moisture which is being carried along with the airstream in the form of droplets. The amount of moisture in 1,000 ft. of free air on a humid summer's day is about $1\frac{1}{2}$ pints and the earlier this can be removed the better.

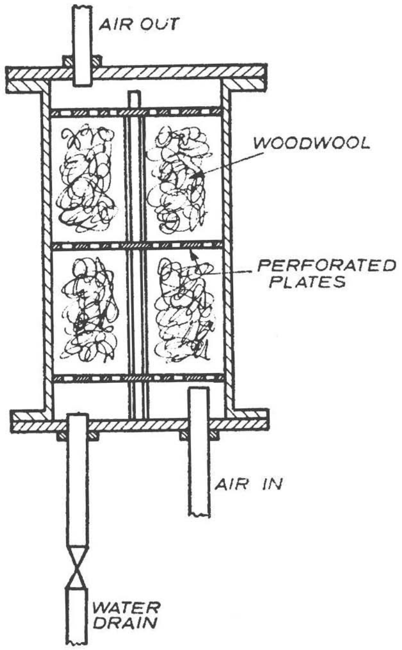

FIG. 11.

Main air filter for fitting immediately after receiver. It should be as large as possible. The woodwool intercepts free moisture which drips down by gravity to the bottom where it is drained off.

A trap such as that shown in Fig. 11 is excellent for this purpose. A pipe of ample dimensions, the larger the better within reason, is packed lightly with a substance such as wood wool, interleaved with occasional perforated separators to prevent it from packing. Trapped water tends to run down and is removed continuously by a drain trap.

The Distribution Main

A typical factory layout as recommended by the B.C.A.S. is shown in Fig. 12. There is principal ring main and a subsidiary ring main. All take-off points are from the top of the main pipes. All pipes slope downwards where possible and drains are fitted at the lowest points. These should be fitted with self-acting traps so that they do not require attention. It will also be noticed that filters and lubricators are fitted in some instances immediately before the tools.

Small Air Filters

It is now common practice to fit a separate filter to each machine or group of machines. To be most effective the filter must be immediately after the tapping from the main line. The air will usually be at ambient temperature and line pressure. After the air leaves the filter it should have no surplus water but will still be at 100% humidity and the flow conditions will inevitably cause the pressure to drop so that the relative humidity may fall to something less than 100% at the point of use. If there is a reducing valve the humidity drops still further and the air would be capable of re-absorbing water should any be present.

In suitable cases an oil mist lubricator can be fitted after the filter and reducing valve, if one is fitted. This draws a small amount of oil into the airstream which lubricates the working parts of valves and cylinders.

A great advantage of adding oil at this point is that it cannot form an emulsion with deposited water, as oil carried over from the compressor tends to do. The provision of the local filter ensures that the free moisture in the air at the tool is negligible.

Air Filter Construction

A conventional air filter is essentially a water separator and operates on an entirely different principle to a filter for, say, hydraulic oil. With an oil filter the filter element is designed to restrain the passage of unwanted particles by making the pores of the element small enough.

An air filter, on the other hand, takes advantage of the differences in viscosity and density

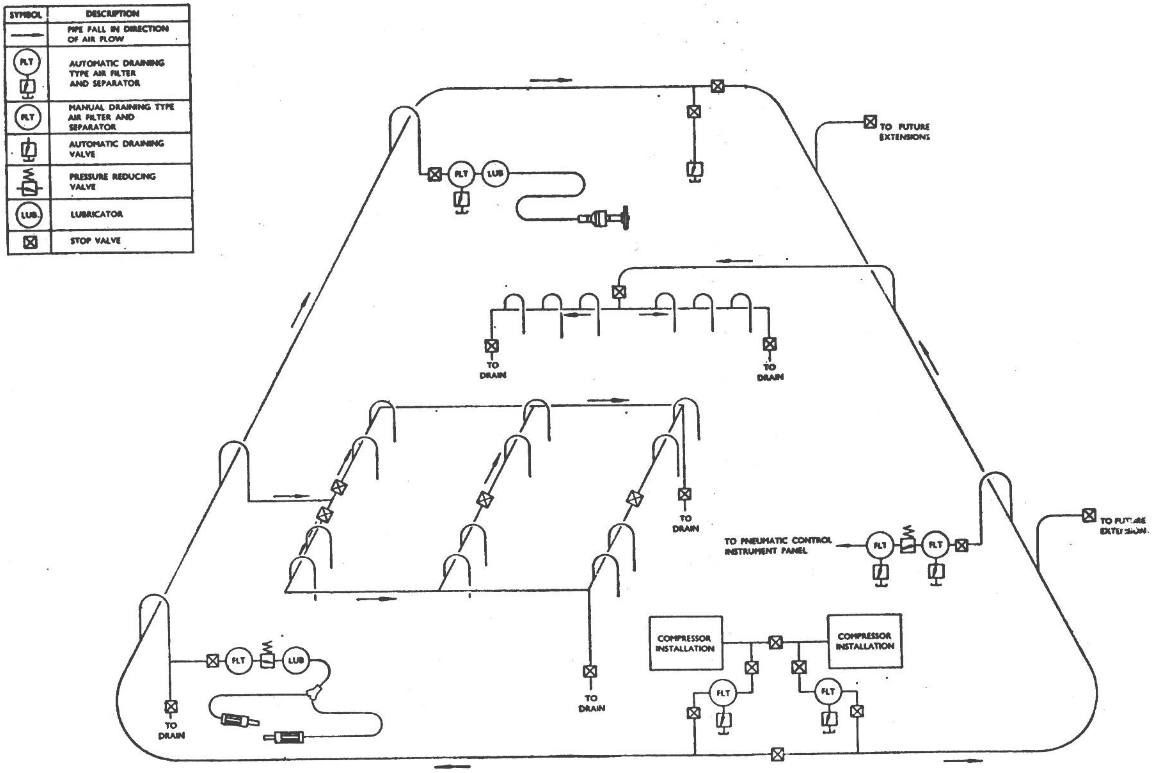

FIG. 12.
Diagram of air distribution network showing main and subsidiary lines and location of filters and oilers.

NOTES
1. Air main drain pipes shown fitted with manual stop valves should for highest efficiency be fitted with automatic draining valves.
2. For maximum efficiency all drains should be automatic. All filters should be of a type equipped with automatic drain.

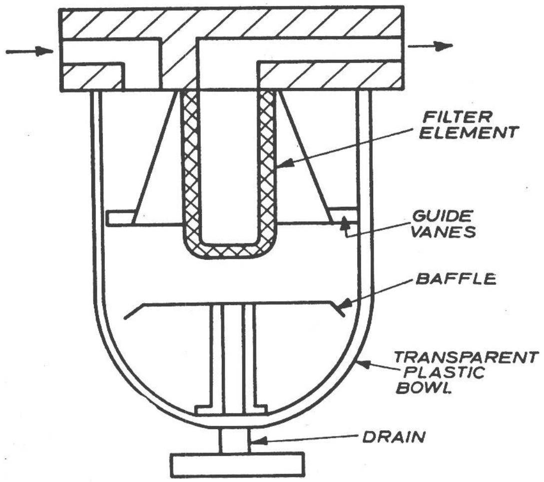

FIG. 13
Conventional bowl filter for filtering at point of use. Air enters from the top, flows round the bowl and then upwards. Water is separated first by the baffle and then by the porous filter, from which it drips down and is collected at the bottom.

of air and water and the filter element is by no means impervious to water.

Probably the most common construction is that shown in Fig. 13 where the air enters at the top and is directed so that it passes downwards along the outside of the bowl. It may be given a circular motion by suitably placed guide vanes.

The water which has to be removed is in a fine suspension in the air, and the idea is to cause the heavier water particles to flow towards the side of the bowl and, as the direction of air flow reverses at the bottom, to remain behind whilst the air passes upwards.

Any water particles which are left are trapped by the tubular filter element in the centre. This element is usually made of porous sintered bronze which forms, in effect, an enormous number of labyrinth passages. These separate out the water whilst allowing the air to pass and the water drains down and falls off into the bowl below. It will be realised that the air at this stage is still at 100% humidity and can only pick up or hold further water if it is moving at a sufficient velocity to hold it against the gravitional force.

A rather different approach is used in the filter shown in Fig. 14 which is intended for comparatively small flows. Here both air and water

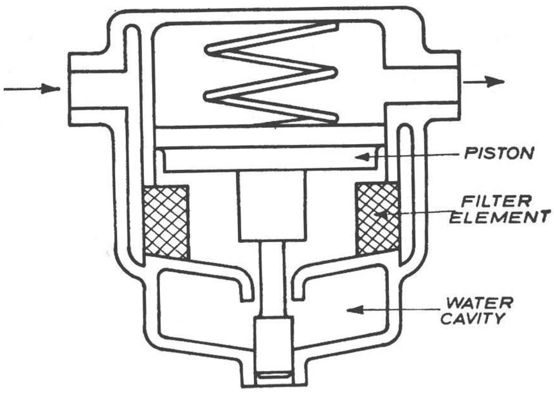

FIG. 15.
Self-emptying filter.

pass through the filter element. The element is made of wool which absorbs the water. The water gravitates to the bottom as the element becomes wetter and drips off into the collecting bowl.

Both the filters just described need emptying periodically, a task which can be overlooked. For large filters a float operated valve similar to a steam trap will do this automatically but is seldom worthwhile on the smaller filters. A self-emptying air filter is shown in Fig. 15. This has a light piston at the lower end of which is a cylindrical plug and which is biased downwards by a light spring. Air enters the piston cavity through the filter element which extracts the water from the air which soaks through the element as it becomes sodden and falls to the water cavity. When a valve is opened the flow of air lifts the piston which closes the upper passage to the water cavity and opens the hole at the bottom, so allowing any accumulated water to be blown out.

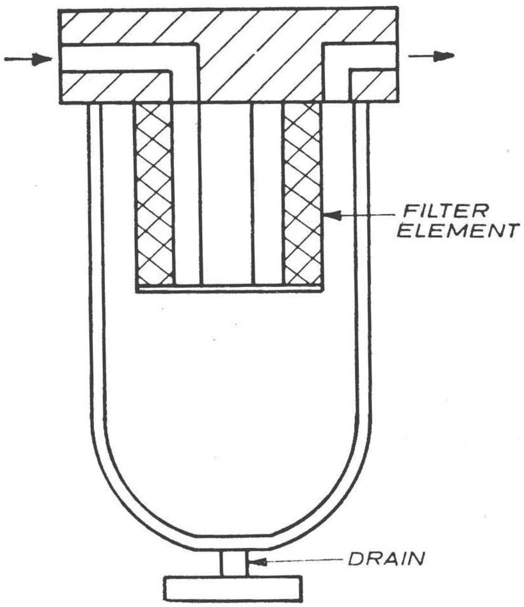

FIG. 14.
Small bowl filter with wool element which absorbs the water and also traps any dirt. The water drips down as the element becomes saturated.

Drying Air

For some purposes air which is dryer than can be produced by mechanical filtration is essential, and it is then necessary to reduce the dewpoint to the required figure. To reduce the moisture content below 100% humidity is relatively expensive. One method is to refrigerate the air to a temperature comparable with the moisture content required, but this is rarely done today.

The most popular method is to use an absorbent such as silicā gel, of which a single pound has an absorbing area of 2,000,000 sq. ft. There are usually two absorbers through which the air is passed alternately. Both absorbers are equipped with heaters and the one not in use is reactivated by heating, a small flow of air carrying the moisture off and discharging it to atmosphere.

The heat required to reactivate the gel is the sum of the latent heat of the evaporated water and the heat losses to atmosphere at the working temperature of about 300°F. An efficient absorber will produce air with a dewpoint of about −78°F, i.e. the air would have to be cooled to that temperature before depositing any moisture.

Checking for Leaks

Leakage is the bug bear of all factory air systems, especially those which have been installed for some time. The power lost through even small leaks can be considerable, as can be seen from the table.

ORIFICE SIZE (INCHES)

	$\frac{1}{64}$	$\frac{1}{32}$	$\frac{1}{16}$	$\frac{1}{8}$	$\frac{1}{4}$
F.A.D. ..	·3	1·2	5·0	20	75
Equivalent B.H.P. (at compressor)	·05	·25	1·25	5	20

(Loss through orifices from 100 psi).

To ascertain what proportion of the air being compressed is lost by leakage is comparatively simple if it is possible to choose a time when the only air being consumed is that going to waste.

The compressor plant is started up in the normal way when the pressure at the compressors is about 10 psi below the normal working pressure. The time taken from the working pressure to be attained is taken. The compressors are then stopped and the time taken to fall to the first pressure.

As the air has been leaking away for the whole time i.e. $T^1 + T^2$ and the compressor has only been working for part of the time T^1, then the proportion of the compressor capacity being lost by leakage is given by $= \frac{T_1}{T_1+T_2}$

This test is so simple to carry out that it should be done at regular intervals. In practice it will probably be found that it is better to work on a pressure range equivalent to some convenient charging time, say two minutes. If this method should be inconvenient then the actual leakage can be obtained from a knowledge of the capacity of the system when the time for the pressure to fall a convenient amount, say from 80 psig to 40 psig is then taken. The approximate leakage rate can then be determined from:—

$$\text{F.A.D.} = \frac{V(80-40)}{t \times 14{\cdot}7} \text{ cu. ft./min.}$$

V = volume of receiver and pipework (cu. ft.)
t = time in minutes.

To take an example from a small installation:—

Volume of receiver 2 ft. dia. × 3 ft. long .. .. = 7·5 cu. ft.
Volume of 200 ft. 2 in. main = 4·3 cu. ft.
Volume of 200 ft. $\frac{3}{4}$ in. branches = ·75 cu. ft.
Volume of remaining lines, say = ·45 cu. ft.

13 cu. ft.

Time for pressure to drop, by measurement .. .. = 15 minutes.

$$\text{F.A.D.} = \frac{13 \times 40}{15 \times 14{\cdot}7} = 2{\cdot}35 \text{ cu. ft./min.}$$

Detection of leaks

Systematic leak detection is best tackled when the pipe work is installed by dividing it into sections by full-way valves and fitting each

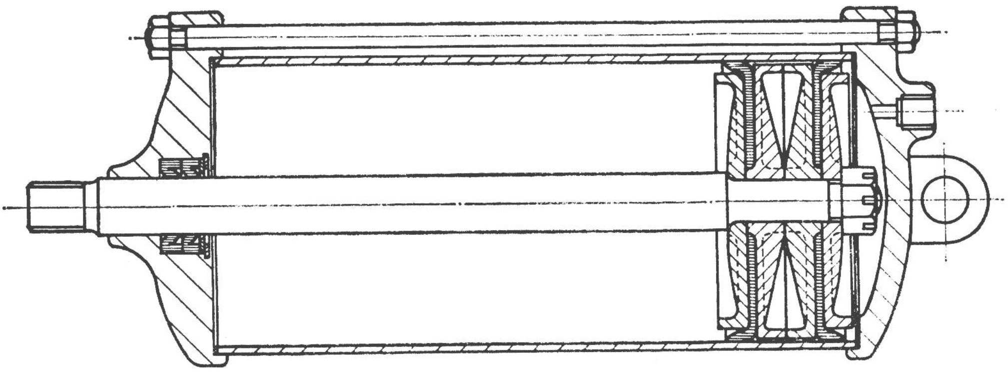

FIG. 16.
Typical double-acting pneumatic cylinder with pivot mounting. On smaller sizes the cylinder tubing may be screwed into the ends, or the cylinder made of extruded light alloy. The latter may be square externally with the ends inserted.

section with a pressure gauge. By isolating each section in turn and watching the rate of pressure drop in the pressure gauge, the worst sections can be located easily. Detecting individual leaks is then usually a question of listening for escaping air when the machinery is stopped and things are fairly quiet.

PNEUMATIC CYLINDERS

The principal method of applying compressed air to automation and the power operation of numerous devices is the cylinder (Fig. 16)—the generally accepted term for the pneumatic linear actuator. A typical cylinder is composed of the barrel, end covers, piston and piston rod together with the appropriate packings.

A cylinder converts the energy in the compressed air into a thrust or a pull. The force is available throughout the stroke of the piston and is the product of the air pressure and the piston area.

Cylinders can be single or double acting. Unlike a hydraulic cylinder the area of the rod, which in an air cylinder does not have to transmit such a large load in proportion to the area, can usually be neglected in calculating the effective area of a cylinder. If pressure is applied to both sides simultaneously it is probable that the piston will not move, being held in position against any differential force by the packing friction.

The pneumatic cylinder is a most useful piece of apparatus for the machine and tool designer as it enables straight line motion to be applied exactly where required, without any mechanical complications such as gears, shafts, cams, etc. The force can be controlled exactly by varying the pressure, whilst the speed can also be controlled with sufficient accuracy for many purposes. It is difficult, however, to stop the piston at a particular point in the stroke other than by obstructing its movement mechanically.

When air is admitted to a cylinder the behaviour of the piston—i.e. its acceleration and velocity—depends on the nature of the resistance to the piston movement and the pressure and velocity at which the air can enter the cylinder.

At one extreme is the condition when the cylinder is connected to a reservoir of air through a large port valve. Opening the valve will cause the piston to be impelled forward at a very high speed, with the risk of damage when it hits the end cover or other resistance.

If there is a resistance to be overcome, say a weight has to be lifted (Fig. 17) the behaviour will be very different. Before the piston can move the pressure must build up to a value sufficient

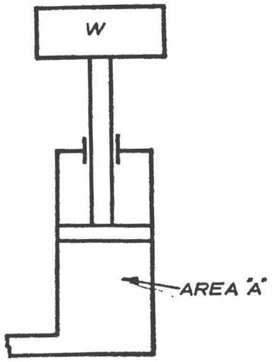

FIG. 17.
When lifting a weight pneumatically the pressure must be rather more than $\frac{W}{A}$ psi. If the weight is supported by compressed air, with the valve closed, it can be raised or lowered a short distance by hand.

both to support the weight— $\frac{W}{A}$ psi and overcome the packing friction, which can be assumed to be about 10% of the load. When the piston starts to move the rate of flow of air will depend on the difference of pressure between the source of supply and in the cylinder, and the resistance to flow in the pipes and valves. If the valve is closed before the end of the stroke, then the piston will continue to move forward as the air expands, assisted by the inertia of the weight. If the cylinder is long enough the piston will stop moving before it hits the cylinder end and will then bounce up and down until friction damps out the oscillation.

When a cylinder is used for raising a weight, as in a hoist, the admission of air will normally be controlled by a hand-operated valve and the speed will be adjusted by the user. When the piston stops with the valve closed, the pressure will be just sufficient to support the weight. It is then possible to raise or lower the load by hand, as might be required when manipulating a heavy part during assembly, without using the valve. The amount of movement possible without excessive effort depends on the volume of air trapped under the piston. The larger the volume, the less the pressure change and therefore the less effort required for a given movement. Packing friction helps to hold the load in a given position and prevents springiness.

A more common arrangement in practice is moving a weight horizontally, as in a machine slide. Here the resistance is largely frictional and the difference between static and dynamic friction is a dominant factor.