TRONDHEIM, Norway 22.-24. JUNE 1982
The Second International Symposium on
WAVE ENERGY UTILIZATION

THE USE OF GYROS AS A REFERENCE FRAME IN WAVE ENERGY CONVERTERS

S.H. Salter

Mechanical Engineering
University of Edinburgh
Scotland

ABSTRACT

Reference frames are a critical part of any wave energy converter. Despite apparent complexity gyros have many attractive features and indeed have a long record of successful duty at sea in the similar task of stabilising ship roll. Characteristics, practical details and reliability are discussed. Official estimates show that the cost of duck-generated electricity from the 1981 design is very close to that of more conventional methods. Reliability is of crucial importance and improvements to assumed levels of reliability could close the remaining gap.

Reference frames

It is quite easy to design an object which will move vigorously in response to sea waves. Indeed it is quite difficult to design one which will not. But if this movement is ever to perform useful work then the designer must provide a reference frame against which his power conversion mechanism can react. The study of such reference frames provides a useful classification for the rich variety of designs of wave energy device. Reference frames are more important and more expensive than the primary interface to the sea.

Although the first duck models were tested as single units in a narrow tank on a rigid axis, we always believed that economic rigid attachments could never survive in extreme waves. Furthermore we wanted to be in deep water and to have a very low dependence on the geology and topography of the sea bed. From the very beginning we planned to use long, round, crest-spanning spines to support a large number of ducks. The idea was that if the spine was long enough to span a variety of phases of wave the net force on the whole structure would tend towards zero. A duck moving up would work against another moving down, somewhere to left or right. The advantages were that we could make the fullest use of available sea front, that we could accumulate power to large amounts, that we shared mooring duty, that we avoided side-to-side collisions, that attachments to waves of opposite phase were achieved with a mechanism which was itself a useful power generator and that nearly all the work could be done in factories on land rather than in the open sea. We recognised from the outset the disadvantages that the stresses in the spine would be particularly severe.

We expected that the power mechanism would be some form of pump acting between duck and spine. We received many congratulations on the simplicity of the pump mechanism proposed in 1973 (1) which is reproduced in Fig. 1.

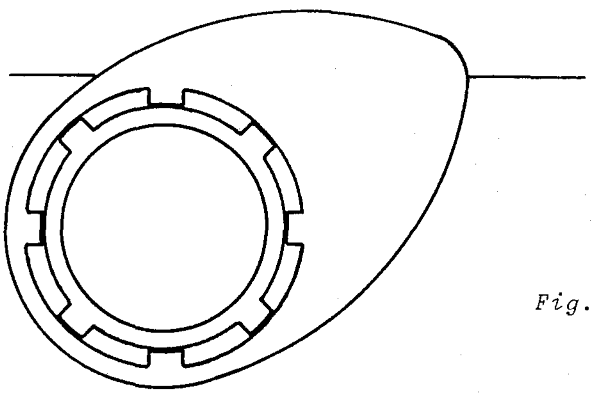

Fig. 1. The delightfully simple first idea.

We knew that these pumps would produce large torsions in the spine. But even worse were the bending moments which were certain to occur in extreme waves. Duck evolution has been dominated by obsessions with bending moments. Everything must be done to understand the statistics of their magnitudes and phases. We had to do all we could to reduce the requirement for strength. We believed that it was wrong to suffer any stress which was not performing useful work or to have any part with a stress exceeding the value it would have at the economic power limit. We had to establish the values of the economic limits and to devise mechanisms to yield to stresses above them. The background to the discovery that the spines did not need to be rigid as well as strong and that they could be completely free in the vertical direction is described in the proceedings of the Gothenburg conference (2). Our current tests on a very long spine model with correctly scaled compliance are confirming the hunch that in random seas bending moments *reduce* above a critical minimum length.

It did not take very long to fall out of love with the simple pumps of fig. 1. We were to discover that considerable performance improvements could be achieved by changing the phase of the power take-off force so that it provided the reactive components of negative spring and perhaps also negative inertia. (3) This meant that we wanted a power conversion mechanism that could be continuously controlled. The 'simple' pump suffered from end stops which would be unacceptable in rough conditions. The end stops of any wave energy mechanism have to be designed with extreme care. Tests in big waves made us realise that it was essential to establish precise design limits for torques, powers and pressures and that it was very desirable to use mechanisms which could, if necessary, be disengaged. We were also learning of the unspeakable horrors of bio-fouling and of the consequent importance of isolating everything from sea water. Finally, we realised that the economics of erratic energy sources would be much better if the device could store energy close to the front of the power train.

It turned out that the use of gyros helped to solve problems in every single area of difficulty.

Gyros

I have never known any idea which received as much initial resistance as the suggestion to use gyros for the power conversion reference frame. But I am happy to report that this initial reaction is almost invariably replaced by enthusiasm in people prepared to understand the technical arguments.

There seem to be several reasons for antagonism. One is a suspicion of the counter-intuitive behaviour of gyros combined with the feelings of intellectual inferiority that they induce. There is also the association with the superb precision demanded from navigation gyros and the thought that this is out of place at sea. Finally there is the need to think of rotations in three dimensions which is much harder than translations in two dimensions. It is quite difficult to write a clear explanation of why gyros behave as they do. I propose to evade the problem by referring readers to better authors (4) (5). Provided we know *what* they do we can use them without knowing why they do it. After all naval architects have been doing very well without understanding the cause of gravity which is fundamental to all their calculations of buoyancy. Electronic engineers can make the cleverest circuits without knowing what electrons are made of.

Gyro characteristics

Let us forget for the time being everything about moments of inertia and angular momentum. Let us treat the gyro and its mounting frames as a sealed black box with two shafts as shown in fig. 2. Let us forget that familiar objects move in the direction of the forces applied to them and that most shafts rotate in the direction of applied torque. Shafts emerging from black boxes do not have to be so boring. They can follow any relationship chosen by the box designer.

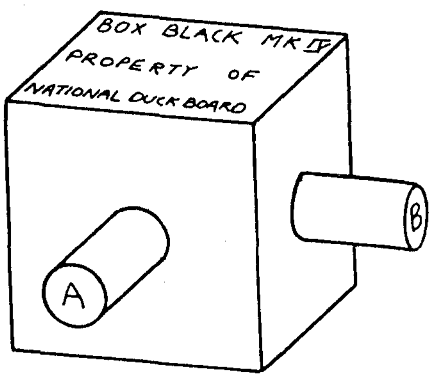

Fig. 2. *An invaluable piece of equipment for every inventor*

The relationship that we specify for our box is that the angular velocity of one shaft should be proportional to the torque on the other and vice versa. The algebra describing this is not daunting:

$$\tau_A \propto \omega_B$$

$$\tau_B \propto \omega_A$$

Proportionalities can be converted to equalities by the use of a constant.

We are free to choose any symbol for the constant and I choose to use the rather odd symbol $I\Omega$. The ◯ keeps the symbols from getting separated. Our equations are not much more difficult than the earlier proportionalities:

$$\tau_A = I\Omega\, \omega_B$$

$$\tau_B = I\Omega\, \omega_A$$

Nobody has said anything yet about input or output; the relationships are entirely symmetrical. Furthermore nobody has said anything about the direction of causality. It is natural to suppose that forces and torques are the causes of velocities and accelerations but there are many cases where accelerations and velocities cause forces. It is best to postpone preconceptions about what causes what. Let us make sure that we understand the black box by considering a few cases of its behaviour when various objects are connected to the shafts.

Firstly we remove all friction springs, dampers and inertial weights from shaft B so that it is totally free to move. This means that τ_B is zero. It follows that ω_A is also zero i.e. shaft A cannot move no matter how large the torque applied to it. Freedom on one shaft has locked the other, and of course locking one shaft would completely free the other.

Let us now connect shaft B to a damping mechanism. Dampers are devices which provide a torque in proportion to and in phase with the applied angular velocity. The damper we connect has damping coefficient K

$$\tau_B = K\ \omega_B$$

Where we have a torque we can replace it with an angular velocity of the other shaft.

Let us change both sides

$$\omega_A I\Omega = K \frac{\tau_A}{I\Omega}$$

or

$$\frac{\tau_A}{\omega_A} = \frac{I\Omega^2}{K}$$

This means that shaft A will feel as if it is connected to a damping device with a value dependent on the reciprocal of the damper on B. Low damping coefficients on one side look like high ones on the other. This seems to be a painless way of obtaining the high damping coefficients which wave energy devices need.

If we double the damping on shaft B, say by adding a second damper in parallel with the first, we halve the apparent damping seen on side A. This means that components which are parallel on one side are seen in series on the other.

Next let us try a spring on shaft B.

Springs are devices which provide forces proportional to displacement. Angular springs produce torques proportional to angle. We get angles by integrating velocities. If the spring constant is S

$$\tau_B = S\int \omega_b \, dt.$$

What does this feel like at shaft A?

$$(I\Omega)\,\omega_A = S\int \frac{\tau_A}{(I\Omega)}\, dt$$

$$\int \tau_A \, dt = \frac{(I\Omega)^2}{S}\,\omega_A$$

This means that the velocity of shaft A will be the integral of the torque applied to it. This means it must feel like an inertia. Furthermore a very small spring at B feels like a very large inertia at A. A spring of zero value would feel like an infinite inertia.

It is not a large step to argue backwards that inertias look like springs and that very small inertias look like very stiff springs.

When we let shaft B be free we were in effect connecting it to inertia, spring and damping with zero value. The rules of the black box make shaft A feel that it has inertia spring and damping of infinite value. It is not surprising that this combination made the shaft immovable.

When one is dealing with mysterious black boxes it is always worth checking that they do not violate the principle of the conservation of energy. Input power to shaft A will be

$$\tau_A \quad \omega_A$$

This is seen to be

$$I\Omega \quad \omega_B \quad \frac{\tau_B}{(I\Omega)}$$

$$= \omega_B \, \tau_B$$

So what goes in to A comes out safely from B and the decencies of energy accounting are given the respect they deserve.

Let us now turn to the innards of the black box. There are no prizes for guessing that the constant $(I\Omega)$, which we used to convert our proportionalities into equations, is made up of the product of the moment of inertia of the Gyro disc abouts its axis ($\frac{\pi D^4}{32} \rho L$ in kgm^2) and its spin in radians per second. It is very convenient that we can vary the 'gyroness' by changing spin speed.

The only thing to spoil the properties of the black box is that the above equations are only strictly accurate when the spin axis of the gyro and the two shafts A and B are mutually perpendicular. When we apply a torque to shaft A the gyro precesses about shaft B to try to align its axis of spin with the axis of the input torque. We have to modify the above equations to take account of this precession. If the angle of precession measured from the mutually perpendicular starting point has the value Φ the exact equation is:

$$\tau_A = I\Omega \; \omega_B \cos\Phi$$

As the angle of precession increases the value of precession velocity also increases.

It turns out that we always want to keep the $\cos\Phi$ bit alongside the $I\Omega$ constant. It seems best to include it in the same box so that the key equation to remember is:

$$\tau_A = I\Omega \cos\Phi \; \omega_B$$

The cosine of 30° is only 13% less than 1 so that quite large angles of precession must occur before the gyro becomes ineffective.

Some gyro mechanisms operate so as to keep all the axes mutually perpendicular. I have not been able to devise one suitable for wave power applications. We seem to be forced into using an arrangement in which the largest wave we can use drives the gyro through its maximum precession angle of Π radians. This introduces a limit to the product of torque and time in addition to the limits to torque and power.

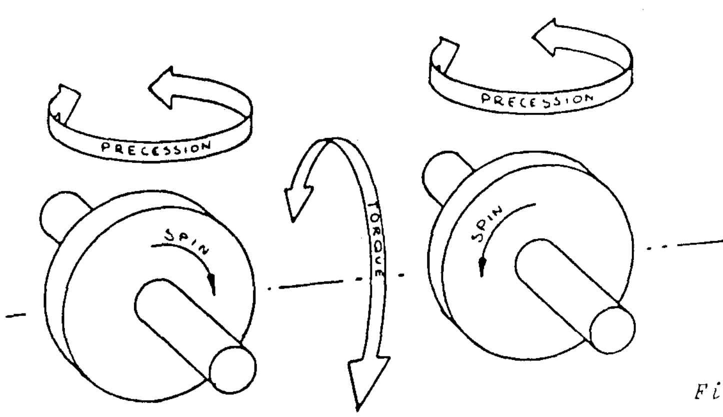

Fig. 3. Gyro behaviour

Let us summarise the properties of the gyro as follows:

Shaft A	Shaft B
Torque	Velocity
τ_A	$I\Omega\cos\Phi \ \omega_B$
ω_A	$\tau_B \div I\Omega\cos\Phi$
Free	Locked
Light damping	Heavy damping
K_A	$K_B = (I\Omega\cos\Phi)^2 \div K_A$
Small spring	Large inertia
S_A	$M_B = (I\Omega\cos\Phi)^2 \div S_A$
Small inertia	Large spring
M_A	$S_B = (I\Omega\cos\Phi)^2 \div M_A$
Series	Parallel
Power = $\tau_A\omega_A$	The same power = $\tau_B\omega_B$

Application to wave energy

There are two ways in which we could have used gyros for the power conversion reference. The simplest one sketched in fig. 3 more fully described in reference (2), had an unforseen snag which we discovered after building an analogue gyro simulator and connecting it to a model in a narrow tank with torques, spins, damping and inertias, all correctly represented. This simple design used pairs of gyros with opposite spin and mounted the power conversion mechanism about the precession axis. We were delighted with the gain (of about three) between the angular velocity of the duck and the precession velocity. This gain could be increased by the easy method of reducing $I\Omega$.

We designed a ring-cam pump around the precession axis. Ring-cam pumps have many attractive features for slow speed hydraulics. The cam shape is shown in Fig. 4.

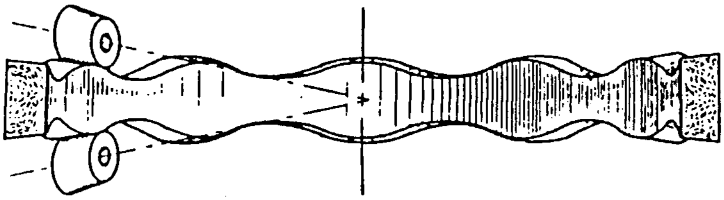

Fig. 4 ring-cam pump

The pumps were fitted with many roller followers and variable displacement was achieved by enabling a varying number of the inlet valves. The pump output fed two high-speed axial piston pump/motors driving each gyro disc and another driving a synchronous electrical generator. The motors driving the gyro disc were programmed to keep the pressure in the feed line from the ring-cams to a chosen value. If, during a burst of wave activity, there was a surplus of energy the excess was used to increase the gyro disc speed by a few rpm. If the ring-cam output was equal to its mean value then the swash plate inside the disc drive-motors moved to the mid-position so that they coasted. But during a lull the swash plates moved to convert the axial piston machines into pumps which drew energy from disc rotation and so maintained full generator output. The gyro drive-motors had also to supply the energy lost in the gyro bearings and as drag exerted by the residual gases around the gyro disc. These losses are a necessary overhead payment for the benefit of gyros. It is after all necessary to use electric light inside conventional power stations, which have 'house loads' of about 5% of full output.

With a stabilised pressure and the very large reserves of energy in the gyros it was not difficult to synchronise the electrical machine. Its drive-motor was programmed to produce a torque corresponding to the required output current. This meant that any amount of power within the generator rating could be drawn for short periods. The only rule was that if the amount taken exceeded the wave input we would gradually slow the gyro spin and that if we drew less, the spin speed would rise. There is an exact parallel with a conventional hydro scheme draining water from a reservoir filled by a river. Indeed we could even import energy from the land and thereby act like a pumped storage scheme.

We claimed the following advantages:

(1) There was no torque reaction on the spine so that the duty required of it was reduced.

(2) There was a useful gain in angular velocity with a corresponding reduction in pump torque.

(3) There was a very large reservoir of stored energy so that constant-voltage, synchronous generation was made possible. This made parallel power collection easy and avoided the losses and expense of rectifiers and inverters.

(4) We could build an oil centrifuge into the gyro shaft and so obtain ultra clean working conditions without the pressure drop of conventional filters.

(5) Most important of all was the fact that we had achieved total isolation of every piece of power-generating equipment from the outside world. All that emerges from the power canister is an electrical cable.

Can any reader see the difficulty?

We had measured the maximum input angular velocities of ducks in a comprehensive set of tests in extreme waves and were able to calculate the corresponding maximum torques that these velocities would have induced. They were high but we thought that they were clearly defined. However it turned out that ducks with the simple gyro arrangement managed to find a way to move *faster* than those with no power take-off at all. It seemed to be wave steepness rather than wave amplitude that was the cause of the problem. The analogue gyro told us that for a very small but nevertheless unacceptable proportion of the time, in seas along the top slope of the scatter diagram, the forces on the gyro bearings were too big because the duck was moving too fast. Furthermore we realised that if for any reason we lost pressure in the precession pump we would make τ_B equal to zero. Gyro action would then require that the duck input velocity would be zero. If the waves did not share this opinion we would have a nearly irresistible force meeting an immovable object. We were failing to danger. The whole emphasis of duck safety design has always been to increase velocity. This seemed now to be in direct contradiction to the requirements of the fast gyro bearings.

The only way out of this impasse was to add a second gymbal frame and place the power conversion mechanism on the front side of the gyro. The gyros are allowed to precess freely thereby locking the rotation of the second frame. The arrangement is sketched in Fig. 5.

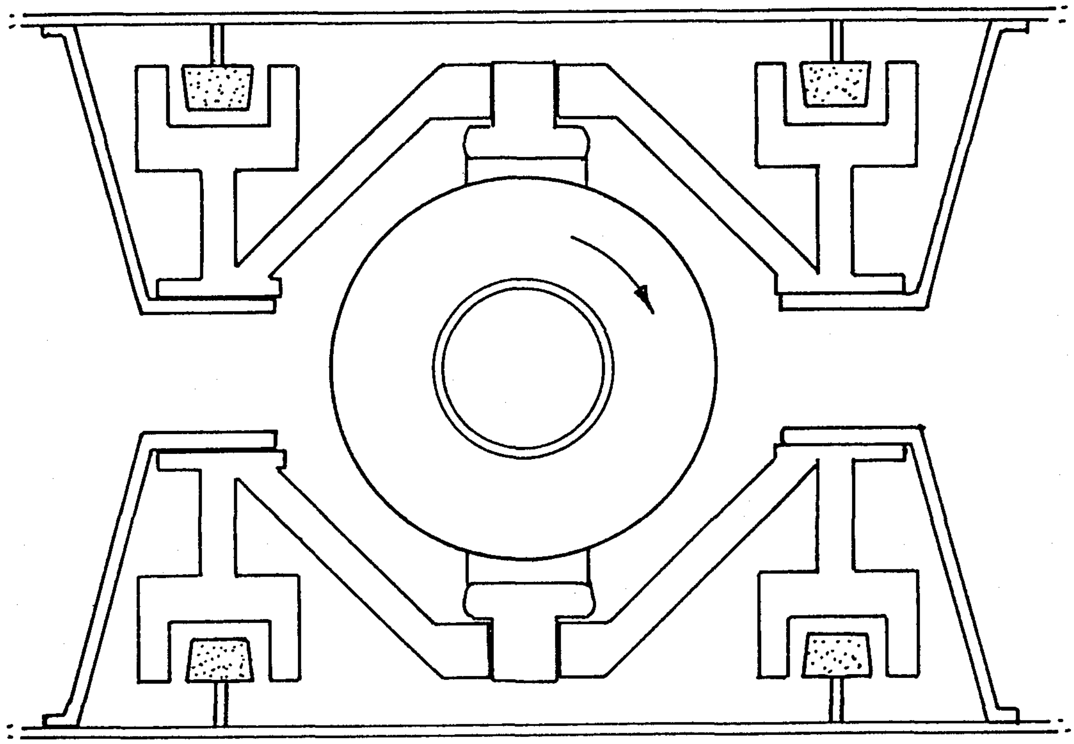

Fig. 5. The second gyro arrangement

We lose the increase in angular velocity and thereby increase the need for input torque. But it turned out that moving the power take-off to the other side of the gyro allows much larger ring-cams to be fitted. The torque developed by a ring-cam pump is proportional to the square of its diameter and so it was easy to provide the extra torque required. The other four advantages were retained. We found that the piping to the gyro drive-motors was more tortuous but that the piping to the generator was shorter. It was very comforting to know that by disenabling the commands to the ring-cam valve-gear we could completely disconnect all incoming power and guarantee the safety of the gyro bearings.

The gyro analogue taught us another valuable lesson. Gyros drift. A centralising spring term can be added to oppose the drift but this spring can set up oscillations. Glen Keller discovered that it is better to use a term he called 'integral of spring'. This is like a spring which pushes back at you harder the longer you have deflected it. It has the same phase as damping and is splendid for stopping drift without starting oscillations.

Gyro Bearings

The full power-conversion torque of the duck is shared between four gyros. It must pass through the slow-speed gymbal frame bearings and the high-speed gyro shaft bearings. The radial forces induced by these torques are large. In addition the bearings have to support the weight of the assembly in whatever attitude the duck may choose. It is interesting to note that forces due to these linear accelerations are much smaller than the useful gyro forces due to angular velocity. The force rating of our 1981 reference design is 170 tons.

The design of the slow bearings was not difficult. We have a guaranteed supply of high-pressure oil and the choice of hydrostatic bearing designs was obvious.

Let me give a brief reminder of the principles. Fig. 6 shows an axial view of a hydrostatic bearing

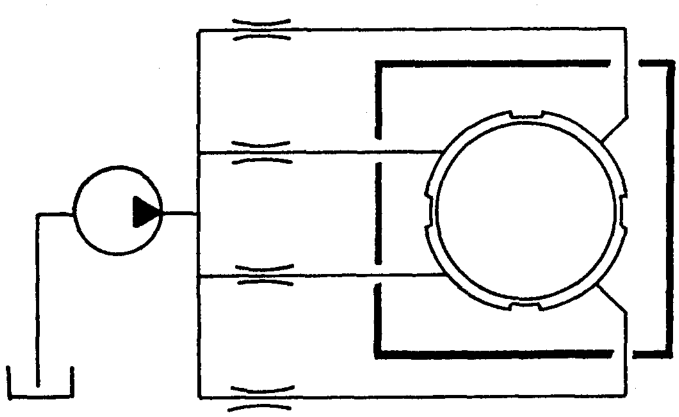

Fig. 6. A conventional hydrostatic bearing

The inner part of the bearing is surrounded by pockets which are supplied with oil through an impedance from a source of high pressure.

If the bearing is unloaded the escape clearances from all pockets and the pressures in all pockets are equal. There is a pressure drop across each impedance so that the pressure in each pocket is about half the supply pressure or perhaps a little more. The supply impedance and the escape clearance are acting like resistors in a potential divider.

If a load is applied to the bearing the central part moves closer to one pocket, closing the escape clearance from that pocket and opening the clearance on the other side. Pressure rises on the loaded side and falls on the unloaded side so that a restoring force is produced to oppose the load.

Hydrostatic bearings can take very high loads. They are more compact than roller bearings and not much more expensive to machine than plain bearings. They have been used for many years for applications such as telescope mountings and advanced machine tools. If they are correctly designed for the loads involved and if they are continuously supplied with clean oil there is never any solid-to-solid contact and their life is indefinitely long. I believe their use would be more common were it not for the trouble of supplying a reliable source of clean high-pressure fluid, which ducks already have.

The design of the fast gyro bearings was a good deal harder. We considered rolling element bearings, active magnetic bearings, conventional hydrodynamic units and finally hydrostatic ones. The rolling bearing design was rejected becuase our calculations showed that its life would be too short. Active magnetic bearings can only develop 'pressures' of about 0.5 N/mm^2 (80 psi) and so we would need to nest several coaxial layers to get enough bearing area. It seemed that the weight and volume would be uncomfortably large. We would also have to provide cooling in the evacuated gyro compartment and the methods of doing this were just as awkward as any of the hydrostatic bearing designs. We therefore rejected magnetic technology but would reconsider it for applications with much higher gyro speeds. The hydrodynamic designs all suffered very high energy losses. We therefore decided to use hydrostatic bearings. But it is not child's play to make a bearing for 170 tons at 1500 rpm which does not have a very large power dissipation. Our design is the work of my colleague Robert Clerk.

There are two sources of energy loss in a hydrostatic bearing. The first is the viscous shear loss in the oil layer between the two moving surfaces. The second is the leakage loss which amounts to the product of the supply pressure from the pump and the volume of its output.

Both sources of loss fall with bearing diameter. We can achieve a small diameter by inverting normal bearing practice and building a stationary spigot inside the rotating gyro shaft as in Fig. 7.

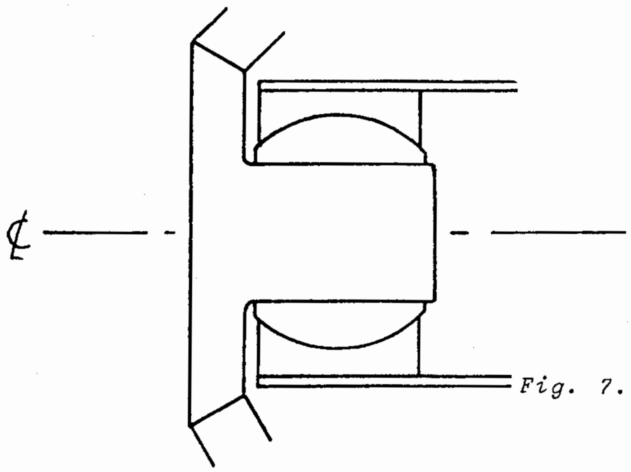

Fig. 7. The inverted configuration reduces diameter

The diameter of the spigot is reduced and its length increased to the point that it is nearly in danger from bending moment failure at the root. It might be thought that bending the central spigot like a banana would jam the fine clearances around it. If the four pockets were equally sized this would be so. But because the forces about one axis are much greater than those about the other we can re-arrange the pockets as in Fig. 8.

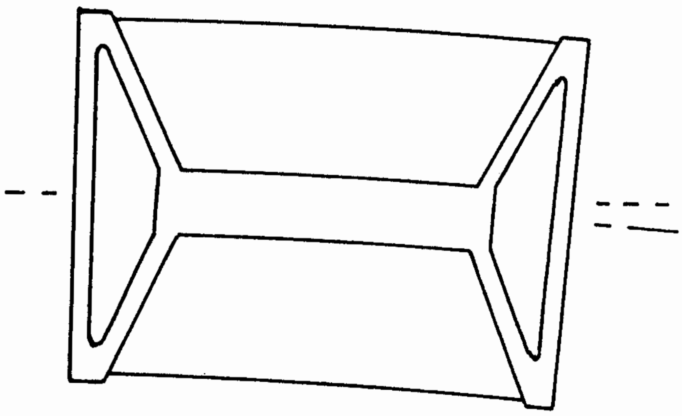

Fig. 8. The new pocket arrangement

This allows 400 microns of deflection at the tip of the spigot without fouling the clearances of only 25 microns.

Most bearings work with highest efficiency when the shear losses and leakage losses are equal. We found that our first selection of land widths, clearances and viscosities gave too much shear and not enough leakage. The introduction of a half-speed sleeve restores the match between the two loss sources.

Finally we grudged the energy lost across the impedance feeding the bearing pockets. Conventional practice uses fixed impedances consisting of capillary paths or orifice plates. If we could make the impedances adjustable we could save much of this loss. The ideal arrangement would be to use high values of impedance in the unloaded condition so that pocket pressures and leakage power are low. When load is applied we would like the impedance on the loaded side to fall so that nearly the full pressure can be supplied to oppose load. The leakage will not be excessive because pocket escape clearances are small. We want the impedance on the unloaded side to rise so that its leakage flow is reduced and no unnecessary pocket pressures are aiding the load.

Fortunately we already have an accurate reliable measure of what the gyro bearing loads will be. They have all come through the slow hydrostatic bearing on the gymbal frame and its pocket pressures can give an exact signal of what is going on. We can use them to drive a sliding spool impedance controlling the fast bearings. This idea allows the fast bearings to use the entire range of their source pressure and achieves spectacular reductions in bearing dissipation. At full load and speed the bearing dissipation is 13.4 kilowatts. Losses steadily reduce at lower power levels and we can arrange for a summer equivalent of hibernation in which essential lubrication services for the whole duck consume less than one kilowatt.

Reliability

The British wave energy programme is subjected to a rigorous professional scrutiny by consultants acting for the Department of Energy. Their successive cost estimates provide an excellent mark of progress. The breakdown of costs highlights the areas in which further work is needed. The estimates of the device teams are not always in exact agreement with those of the consultants and the differences provide fruitful ground for discussions in the following year. The biggest single subject for disagreement in 1981 was the assumption made about the loss of output due to imperfect reliability and the consequent necessary expense of maintenance. Were it not for these two factors the cost of electricity delivered to the grid by the ducks of 1981 would be below the cost to domestic consumers on the UK mainland. It is still too high to compete with the sent out cost of electricity from new coal plant and far above the cost to consumers in Norway. Wait until 1991. Clearly we have to give the question of reliability a very great deal of attention.

It may be helpful to draw up a list of all the sources of short life and bad reliability.

SOURCE OF PROBLEM	Duck	Brand X
Heat, thermal gradients and changes	NO	
Dirt, dust, sludge, grit, ash	NO	
Imprecise knowledge of duty requirement	NO	
Chemical attack, corrosion, oxygen, ozone	NO	
Animal attack, vegetable growths	NO	
Micro-biological activity	NO	
Neutron embrittlement, ultra violet radiation	NO	
Pressure transients	NO	
Absence of lubrication, clogged oil filters	NO	
Heavily loaded solid-to-solid moving contacts	NO	
Incompetent inspection and maintenance	NO	
Salt deposits, ice, snow, frost	NO	
Concrete spalling, rust flakes	NO	
Vibration, acceleration	YES	

We do not believe that equipment at sea is subjected to supernatural emanations induced by the interaction of cosmic rays and water. If something fails, it does so for an identifiable physical reason which can be studied and remedied.

Nearly every unpleasantness is absent from the working conditions of duck power canisters. We enjoy constant temperatures, ultra clean working conditions, exact specifications of the duty and freedom from chemical and biological problems. We have no adiabatic icing difficulties. We are free from the ministrations of fitter's mates. Nearly every bearing uses full hydrostatic lubrication. Where this has not been possible, in particular in the contact between the ring-cam and its rollers, we can compare our requirements with those of rolling bearings and railway wheels. Hertzian stress calculations show that 90% of our rollers will last for a hundred years and we can pension off those that do not. If we compare our poppet valves with those of diesel engines in heavy goods vehicles or the leakage rate of our power canister with that of industrial cryo-flasks, we get similarly encouraging conclusions. Our main stressors are the high frequency vibration from the spin of gyro discs and generator rotors and the low frequency accelerations from wave action. We must balance our rotating machinery sufficiently well to remove this hazard. We can do nothing about the sea motions so duck equipment must stand up to accelerations at the wave spectrum which we estimate at $\pm$ 1.5g. It turns out that the stresses induced in the machinery by these accelerations are far below those associated with useful gyroscopic effects. The only *technical* objection we have so far received from reliability experts is that the nuts will always come loose. We would welcome further technical criticism of a precise, numerical, testable nature from participants at this conference.

Conclusion

Gyros can be used to provide a reference frame for the torque inputs to a wave energy device. They offer advantages in removing the need for torsional strength in long spines and providing hermetic isolation of the entire power canister from the outside world. Internal working conditions are so very good that it seems wrong to apply reliability predictions based on experience with normal exposed marine applications. Cost of the gyro mechanism can be justified by the value of energy storage at the front of the power train and the consequent possibility of synchronous generation and transmission at mean rather than peak levels.

Acknowledgements

Eric Laithwaite taught me about gyros.

Robert Clerk is teaching me about high pressure oil.

Full scale design work is carried out in collaboration with engineers from the Laing organisation and the Scottish Offshore Partnership.

The United Kingdom Department of Energy supports the wave energy team at Edinburgh University.

References

(1) Salter, S.H., 'Wave Power' Nature, vol. 249, pp. 720-724, June 1974.

(2) Salter, S.H., 'Recent Progress on Ducks' 1st Symposium on Wave Energy Utilization, Gothenburg, October 1979

(3) Salter, S.H., Jeffrey, D.C., Taylor, J.R.M., 'The Architecture of Nodding Duck Wave Power Generators'. The Naval Architect pp. 21-24, January 1976.

(4) Inglis, C., 'Applied Mechanics for Engineers', Chap. 20, Cambridge University Press, 1951 (republished Dover 1963).

(5) Arnold, R.N., Maunder, L., 'Gyrodynamics and its Engineering Applications', Academic Press, New York 1961.