PARROTS AND MONKEYS

Introduction

Before this series of experiments, Narrow Tank Duck control strategy
had been researched by varying damping and spring in the three modes,
NOD, HEAVE and SURGE (refs 1,2).   The justification for using this
simple system was that efficiency is most important when the sea is
small and the response of the Duck at its most linear.   By the time
non-linearities are significant, survivability has superseded efficiency
as the prime consideration.   Parrots and Monkeys represent an
experimental approach which can automatically optimize productivity
in whatever way, linear or non-linear, is best.   We describe the
technique and give ways of interpreting the data generated.

How it works

Duck DO027 was mounted on the Pitch-Heave-Surge rig in the Narrow
Tank.   The Duck's width is 76% of the tank, equivalent to the Duck
width/Duck pitch of the full scale spine based design.   The model
is 1/140th of full scale.

Narrow tank mixed sea experiments are now done using a PET computer
to generate a quasi-random sea.   When an experiment is repeated,
the waves are identical and if the control knobs have not been touched,
the time series of the Nod Torque, Heave Force and Surge Force commands
to the Duck are also identical.   We tried recording the following
signals:

             Wave Maker command
             Nod Torque command
             Heave Force command
             Surge Force command

for the duration of an experiment, and replaying them so that the
system was driven by the recorded signals and not by damping and spring
as normal.   The performance was within 1 or 2% of the original.
This technique is what we refer to as Parrotting.

Fig 1 is a block diagram showing signal routing.   Using the 'Zookeeper' one can switch model control from manual to either the PET or the D-A converters run by the SIRIUS.

The recording was done by sampling each signal 20 times a second and storing it in a SIRIUS computer memory.   Initially we used 20 second experiments to validate the system.   This gave us an array of 400 numbers per command - these we refer to as channels.   Clearly if some of these numbers were changed, or Monkeyed with, the performance would be expected to suffer, but this might not always be the case. It is easy to tell a 'good' change, since the productivity increases. During a 'Parrot and Monkey' experiment, the test is re-run several hundred times.   A random change is made to one of the channels every run and if it results in an improvement, it is preserved;   otherwise it is rejected.   By careful choice of the bounds of the random change we can get as many as one in five accepted (see Appendix A).

Fig 2 shows the printout from a typical test.   A line is printed each time there is an improvement;   it tells which channel (N, H or S) was changed and gives information about power, force and waveheight. The improvement is finally measured by doing 5 Parrot runs using the pre-Monkeyed recordings followed by 5 post-Monkeyed and taking their respective averages;   this removes the effect of water level or temperature change during the experiment.   In this printout one can see that improvements were still being made when Monkeying was stopped after 63 runs.

Choice of random change (Monkeying)

A random choice is made between Nod, Heave and Surge, and a random delay of between 7 and 20 seconds is chosen.   (The first 7 seconds are needed for the sea to reach the model.)   The computer has earlier calculated the maximum value in any channel.   The change to be introduced is got by multiplying the appropriate maximum value by a random number whose value lies between the bounds we can set: usually −0.6 to +0.6.   This change is added to the three samples that appear immediately after the delay.

In a typical test run (No. 23D) repeated 74 times without Monkeying, the productivities had a standard deviation of 0.88%. If there are three channels, each of 400 samples, contributing towards the total productivity during a 20 second experiment, and each sample has an equal weighting, the maximum change permitted in one run should only increase or decrease the power by 0.2%. This improvement/noise ratio is the main drawback of the system. In fact it is not as bad as it appears: weightings are not equal, and noise works in our favour if we can do enough runs.

Fig 3 shows what can be achieved by starting with an un-optimized Duck. Over 94 runs we have an improvement of 16% and the graph shows no sign of levelling off.

To speed things up, especially when doing experiments longer than 20 seconds, we tried introducing several changes per run. The 'goodness' of a change was judged by looking at the productivity in the period immediately following it. We investigated the effect of varying this 'power window' width. We measured the total productivity change over a run and then the change in different window widths. Fig 4 shows the results: we can reckon that 95% of a productivity change is present in the first 15 seconds after the Monkey change that produced it. This is unfortunately not a hard and fast rule: it works for small seas but as soon as the Duck is near capsize, a small Monkey change in the right place can affect the whole of the rest of a run.

The bulk of our experiments were done with 4 or 5 changes per 51 second run.

Criteria for acceptance

The simplest way of judging the 'goodness' of a change would be to compare it with the previous run, but for this to work correctly we would only be able to Monkey every second one. Initially what we did was to compare the productivity of a run with the previous 'best productivity'. This worked well where there were big improvements to be had. However, when Monkeying with a near-optimal Duck, long

periods would elapse with no change being accepted: this was to be
expected. But were we rejecting good changes that were not being
helped by noise? To accept more changes, we arranged that the 'best
productivity' fell each run. The rule was that a proportion of the
original 'best productivity' (usually 0.1 or 0.05) was subtracted
and replaced by the same proportion of the productivity from the last
run. (The proportion of the original retained is called 'averaging
time const' and included in the heading to each printout.) We found
that this modification increased the rate of improvement in most cases,
but did allow some experiments to show a net reduction at the end.

Making sense of results

At the end of a Parrot and Monkey experiment we have a set of recordings
- these are:

     a) Nod velocity
     b) Heave velocity
     c) Surge velocity
     d) Nod torque command (Monkeyed)
     e) Heave force command (Monkeyed)
     f) Surge force command (Monkeyed)

The first three present knowledge of the Duck performance that is
always available through measurement channels. To obtain the
improvement that the Monkeys have given us in a real situation, we
need to find a relationship between the first three channels and the
last three.

The sort of relationship one would expect would be of this form:

Nod torque = (K_1 x Nod velocity) + (K_2 x Nod displacement)
+ (K_3 x any other term) + (K_4 etc .....

and similarly for Heave and Surge. In this example K_1 is damping
and K_2 spring. The technique we used for finding these coefficients
is as follows (see Fig 5).

Stage 1   Correlate the recorded command signals with all the channels,
           including new ones got by integrating the velocities and
           any other tricks that come to mind

Stage 2   Select the channel with the strongest correlation and get
           its 'correction coefficient'.   This is defined thus:

$$CC = \frac{\Sigma AB}{\Sigma B^2}$$

           (A = command channel   B = strong correlation channel)

Stage 3   Subtract CC x B from A.   This gives us a new channel, let's
           call it C, which has a zero correlation with B.

Stage 4   Substitute C for the command channel and return to
           Step 1.

Taking the example earlier, if A is the Nod torque command then the
strongest correlation would be with Nod velocity and the CC given
by Step 2 would be the damping coefficient, K_1.   The new channel,
C, would represent Nod torque without the damping term, and repeating
Step 1 would show a strong correlation between C and Nod displacement.
The CC term obtained from these two will be K_2, the spring coefficient.
This procedure can be repeated until the RMS of C is a small fraction
of the original command channel, or C is mainly noise that doesn't
correlate strongly with any channel.

Since this technique was developed it has also proved useful in
interpreting Wide Tank data.

Results

Most of the Parrot and Monkey tests were trials done to debug and
checkout the experimental set up (which now comprises 2 computers,
4 interfaces and a hefty amount of software).   As soon as the system
was working correctly, we were able to try out 'Coefficient Monkeying'.
Initially this means varying the damping and spring coefficients,
rather than the time series;   in the future we intend to add extra

terms to the control strategy.   However, even this simple system rapidly demonstrated that it could obtain greater improvements more quickly than 'Parrots and Monkeys' and we changed over to it.   'Variable Coefficients - Interim Results' contains the improvements obtained in the first 8 of the 46 spectra: the remaining 38 have now been done.

References

1.   Edinburgh Wave Power Project : 2nd Year Report, Section 22.

2.   Edinburgh Wave Power Project : 4th Year Report Vol 1, Section 4.

Appendix A

In the first tests of Parrots and Monkeys the recording was played back without any smoothing - a time series showed the typical staircase waveform, with steps every 50 milliseconds.   We found we could obtain suspiciously large improvements.   It turned out that the Monkeys made the power flow from the Duck oscillate violently, being positive at the instant it was sampled and negative for the rest of the 50 millisecond period.   To avoid this problem, we performed a digital partial differentiation (high pass filter) on the recorded command signals, and fed them to the model through partial analogue integrators (low pass filters).   These two processes cancel each other, leaving us with a smooth waveform.

NARROW TANK BLOCK DIAGRAM
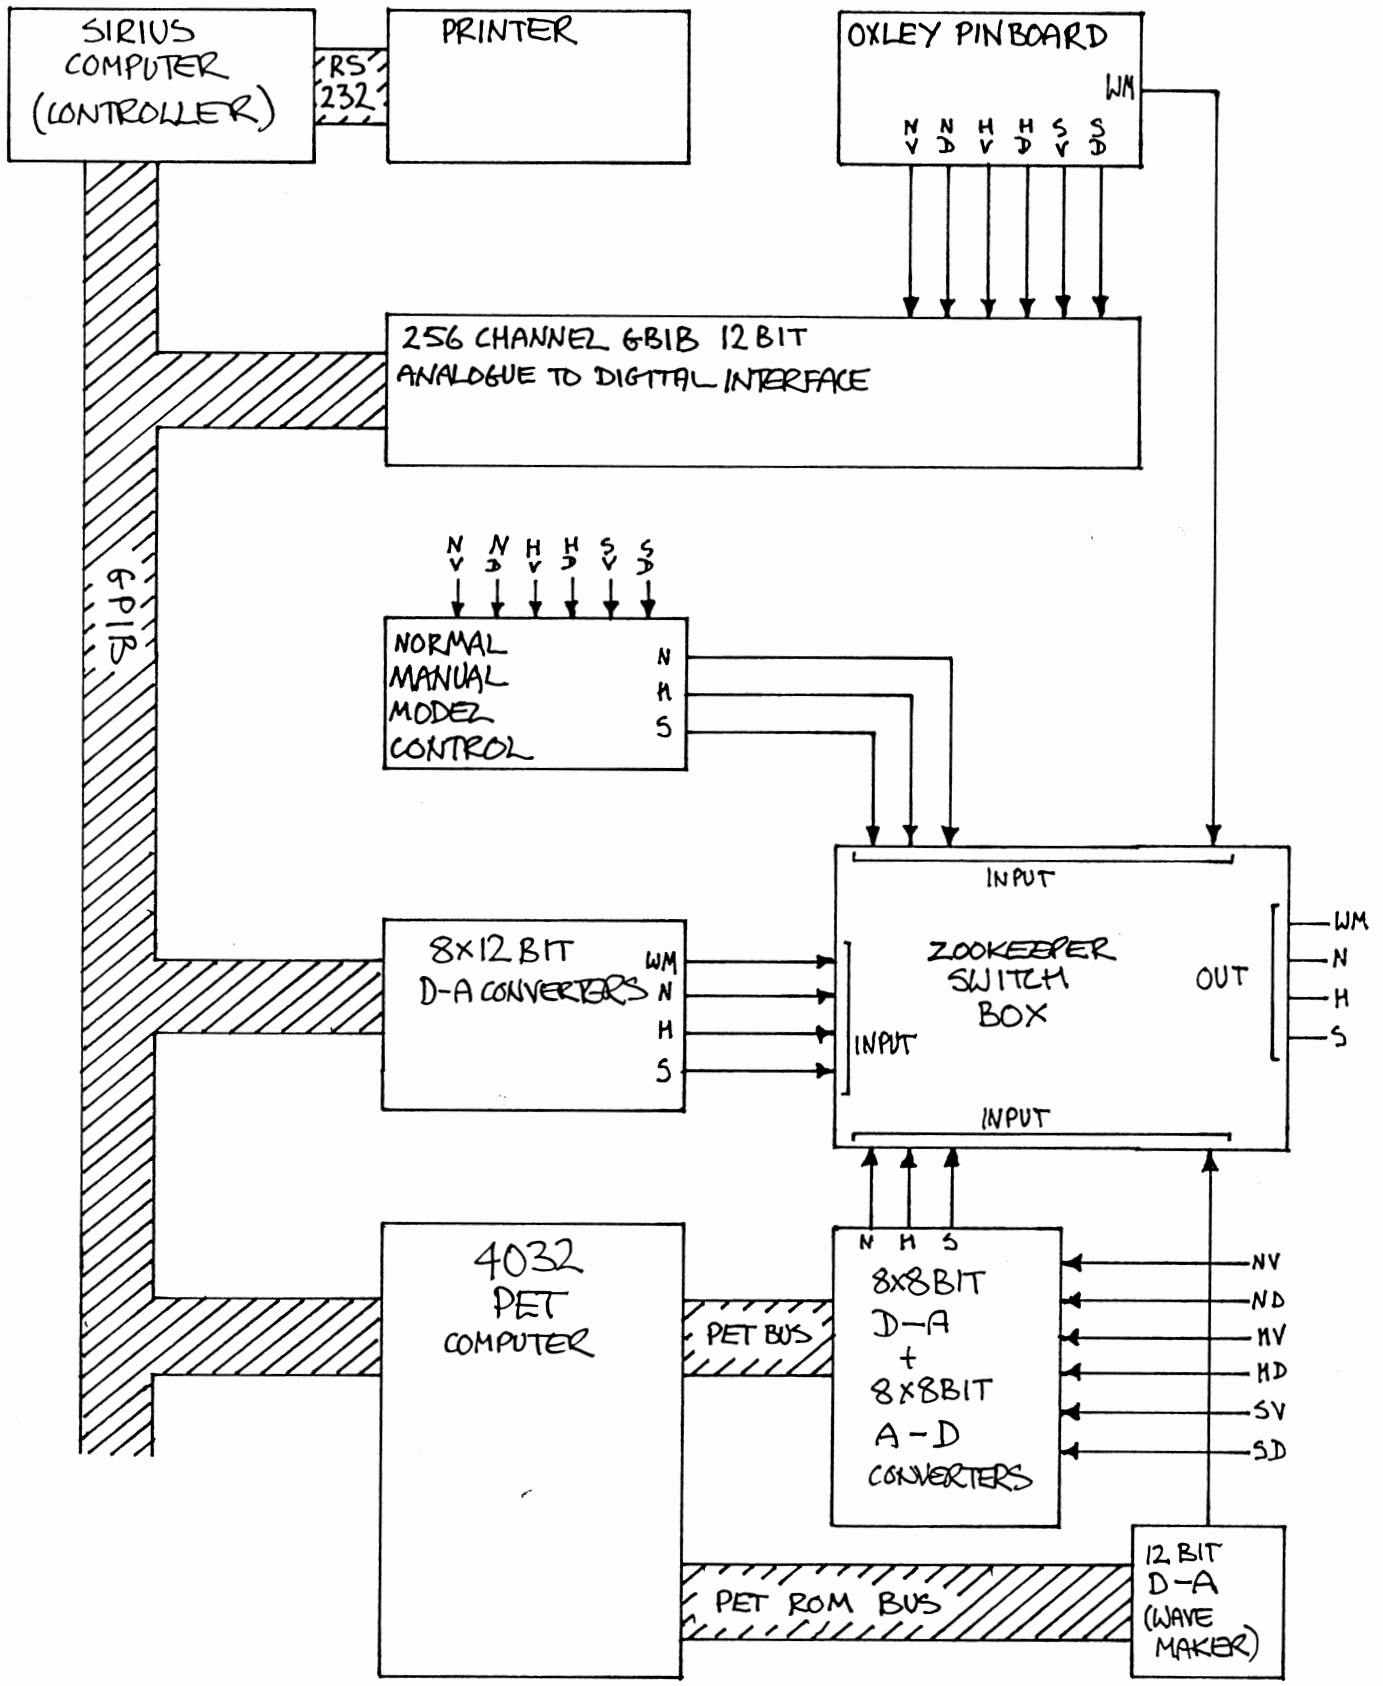
KEY: N = NOD   V = VELOCITY
     H = HEAVE   D = DISPLACEMENT
     S = SURGE   WM = WAVEMAKER

Figure 1

```
                Date......................  29/6/82
                Title/filename...........  TEST#23D
                Sampling rate............  20 Hz
                Sample period............  20 seconds
                Power measured after.....  7 seconds
                Change length............  4 samples
                Random seed..............  4458
                Averaging time const.....  .95
                Max. monkey change.......  .7
                Offsets zeroed every.....  10 runs
                The Intelligent Controller is on line.

   MONKEYING....

Run          mW        mW        mW        mW               Nm       N         N       cm     cm
   1 (S) N 61.5 S 21.2 H  4.1 TOT 86.8 CH  0.0% N0.005 S-0.03 H-0.42 A0.77 B0.49
   6 (H) N 62.0 S 20.9 H  4.5 TOT 87.4 CH  0.6% N0.005 S-0.06 H-0.47 A0.77 B0.49
   7 (N) N 62.2 S 21.4 H  3.9 TOT 87.5 CH  0.8% N0.005 S-0.05 H-0.42 A0.77 B0.49
  11 (N) N 61.2 S 20.6 H  5.6 TOT 87.5 CH  0.8% N0.005 S-0.01 H-0.41 A0.77 B0.49
  26 (N) N 61.5 S 20.2 H  5.4 TOT 87.2 CH  0.5% N0.005 S-0.04 H-0.44 A0.77 B0.49
  35 (N) N 61.8 S 20.5 H  4.8 TOT 87.1 CH  0.3% N0.005 S-0.03 H-0.44 A0.77 B0.49
  39 (H) N 64.4 S 19.0 H  3.6 TOT 87.1 CH  0.3% N0.005 S-0.04 H-0.45 A0.77 B0.49
**************************************************************************************
          We seem to have started pessimising, so I've set TCONST to 1
**************************************************************************************
  40 (H) N 64.9 S 19.9 H  2.5 TOT 87.3 CH  0.5% N0.005 S-0.02 H-0.43 A0.78 B0.49
  43 (H) N 63.4 S 19.6 H  5.4 TOT 88.3 CH  1.8% N0.005 S-0.01 H-0.54 A0.77 B0.49
  45 (H) N 64.2 S 20.1 H  4.2 TOT 88.4 CH  1.9% N0.005 S-0.02 H-0.51 A0.78 B0.49
  50 (S) N 63.8 S 19.7 H  5.0 TOT 88.5 CH  1.9% N0.005 S-0.00 H-0.56 A0.77 B0.49
  51 (N) N 63.8 S 19.9 H  4.9 TOT 88.6 CH  2.1% N0.005 S-0.01 H-0.42 A0.78 B0.49
  54 (H) N 64.5 S 19.2 H  5.3 TOT 89.0 CH  2.5% N0.005 S 0.01 H-0.47 A0.78 B0.49
  63 (S) N 64.9 S 19.1 H  5.3 TOT 89.3 CH  2.9% N0.005 S-0.00 H-0.45 A0.78 B0.49

TEST#23D             PARROTTING ONLY              29/6/82

Run          mW        mW        mW        mW               Nm       N         N       cm     cm
   1 (0) N 65.6 S 19.8 H  2.4 TOT 87.7 CH  0.0% N0.005 S 0.01 H-0.43 A0.78 B0.49
   2 (0) N 65.8 S 19.8 H  2.3 TOT 87.9 CH  0.2% N0.005 S 0.03 H-0.42 A0.78 B0.49
   3 (0) N 66.1 S 19.7 H  1.5 TOT 87.2 CH -0.6% N0.005 S-0.01 H-0.38 A0.78 B0.49
   4 (0) N 64.8 S 19.3 H  4.3 TOT 88.3 CH  0.7% N0.005 S 0.04 H-0.46 A0.77 B0.49
   5 (0) N 65.1 S 19.1 H  4.6 TOT 88.7 CH  1.1% N0.005 S 0.03 H-0.43 A0.78 B0.49
The average is 88.0 mW


TEST#23              PARROTTING ONLY              29/6/82

Run          mW        mW        mW        mW               Nm       N         N       cm     cm
   1 (0) N 60.2 S 22.5 H  1.1 TOT 83.8 CH  0.0% N0.005 S-0.02 H-0.29 A0.78 B0.49
   2 (0) N 60.2 S 22.3 H  1.0 TOT 83.4 CH -0.5% N0.005 S-0.00 H-0.28 A0.78 B0.49
   3 (0) N 60.0 S 21.5 H  3.2 TOT 84.7 CH  1.0% N0.005 S 0.00 H-0.40 A0.78 B0.49
   4 (0) N 60.0 S 22.1 H  3.0 TOT 85.1 CH  1.5% N0.005 S 0.00 H-0.35 A0.78 B0.49
   5 (0) N 60.3 S 21.9 H  2.3 TOT 84.5 CH  0.8% N0.005 S-0.01 H-0.32 A0.78 B0.49
The average is 84.3 mW


****************************************************************************
                    THE IMPROVEMENT IS 4.4%
****************************************************************************



                                                         Figure 2
```

Plot showing productivity increase VS run number.

Data collected 16 Sept. 1981.
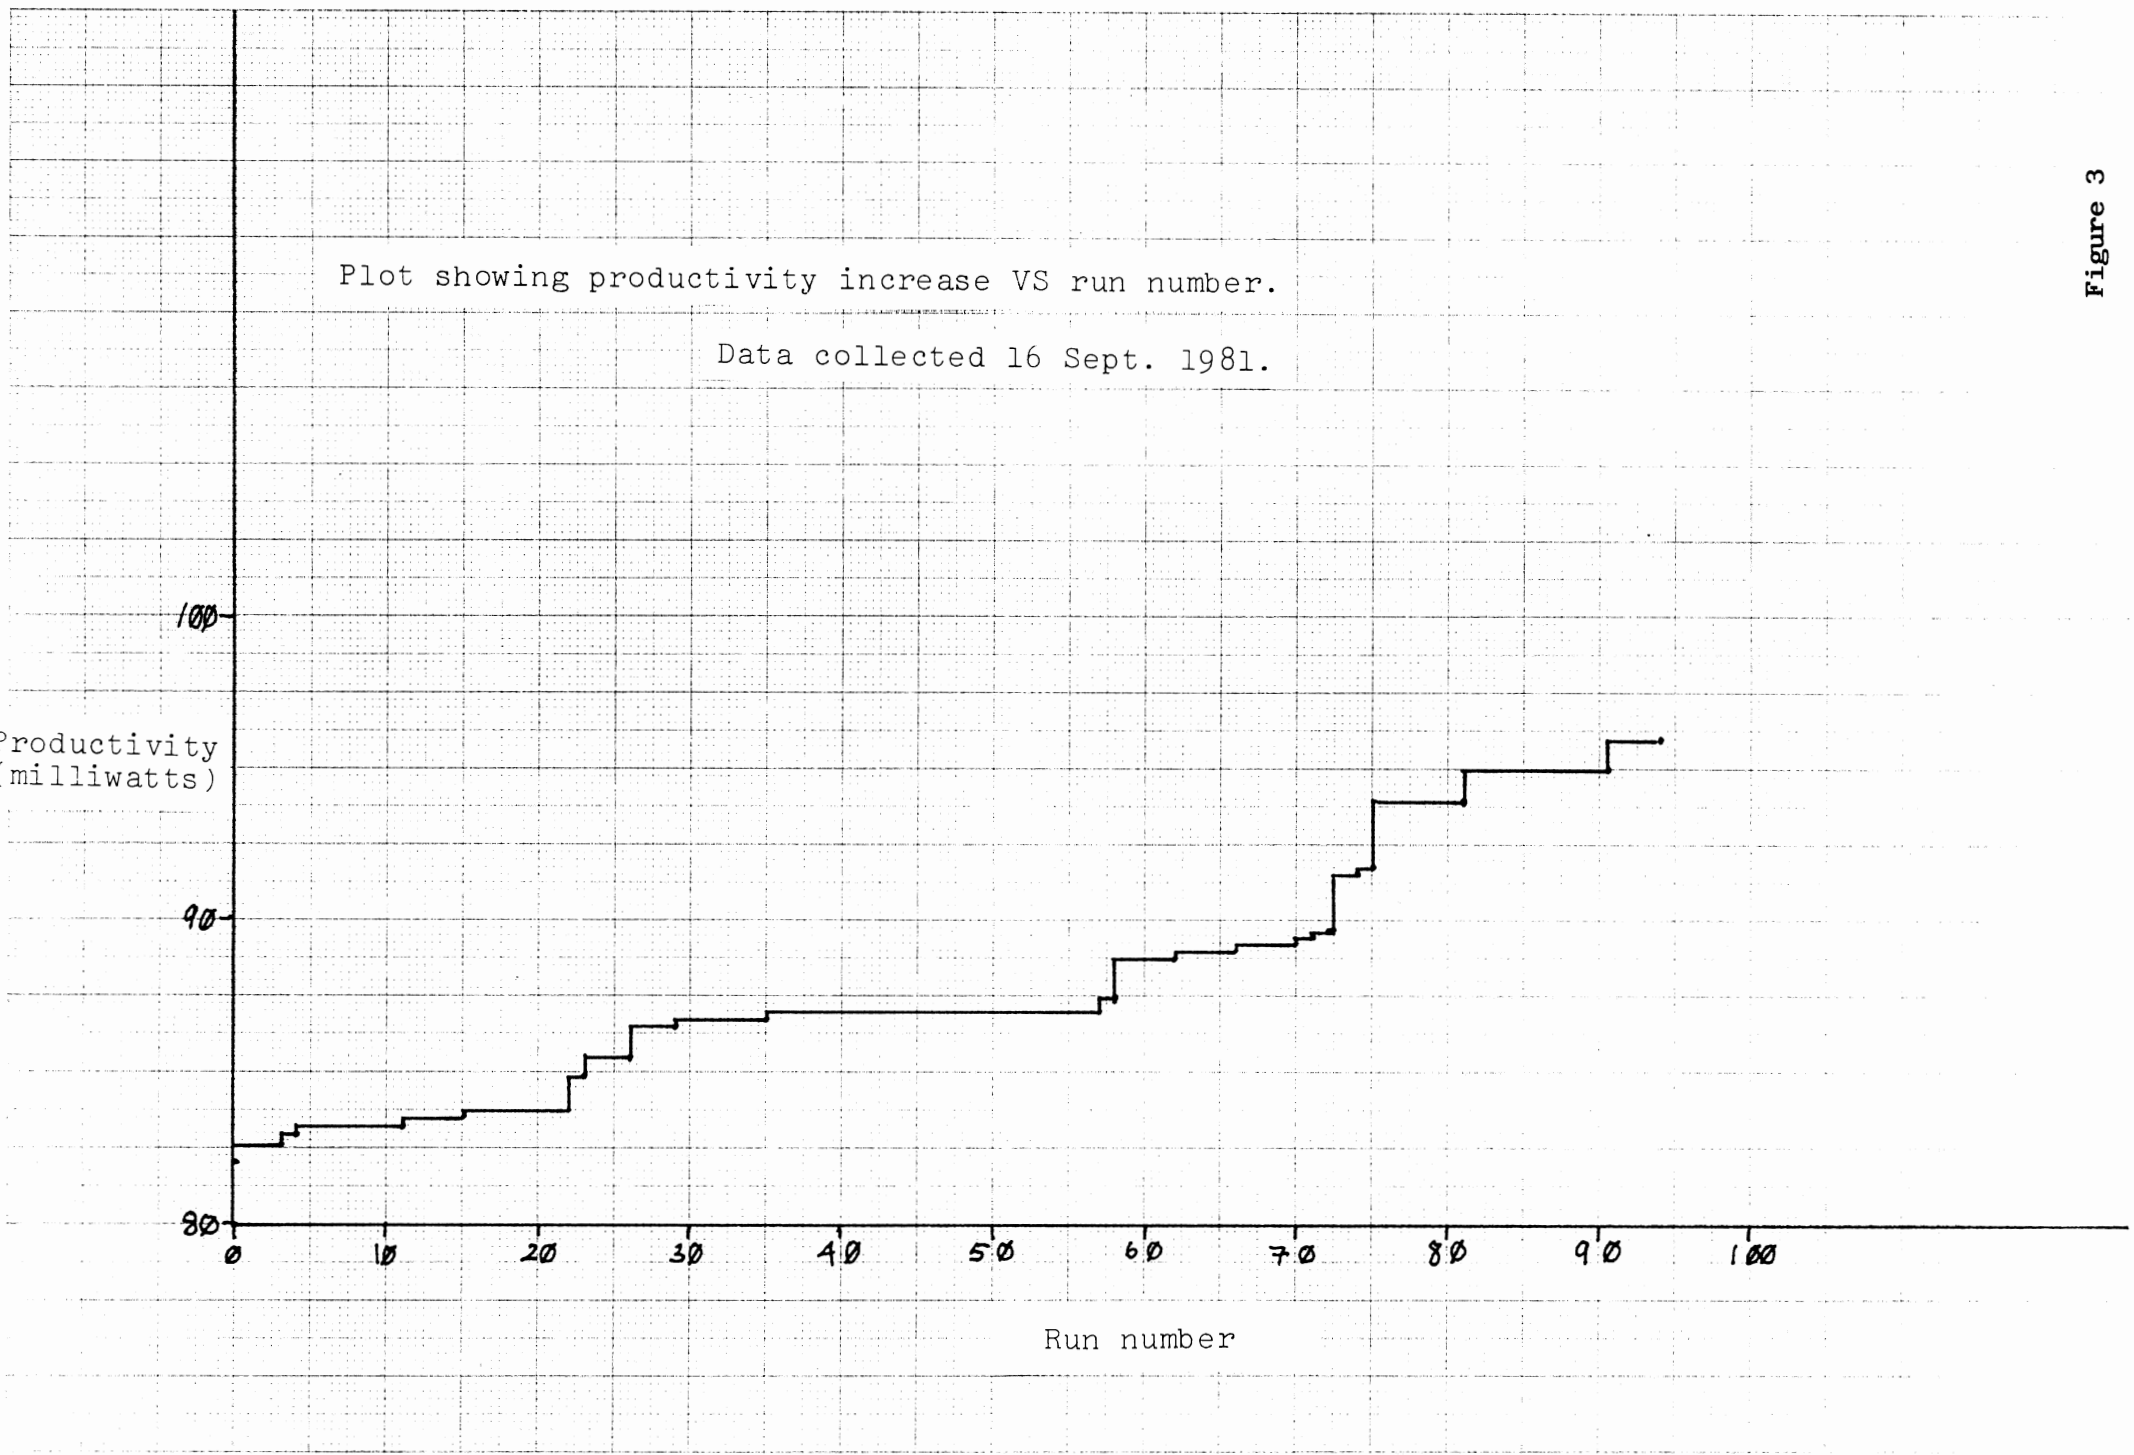
Productivity (milliwatts)

Run number

Figure 3

Plot showing change in productivity during a 'power window' expressed as a percentage of the total change in productivity over a run VS duration of 'power window'.

Time (s)

Figure 4

Nod velocity (rms)   = 0.84 rad/s

Nod displacement (rms) = 0.18 rad

Nod torque   (rms)   = 0.05 Nm
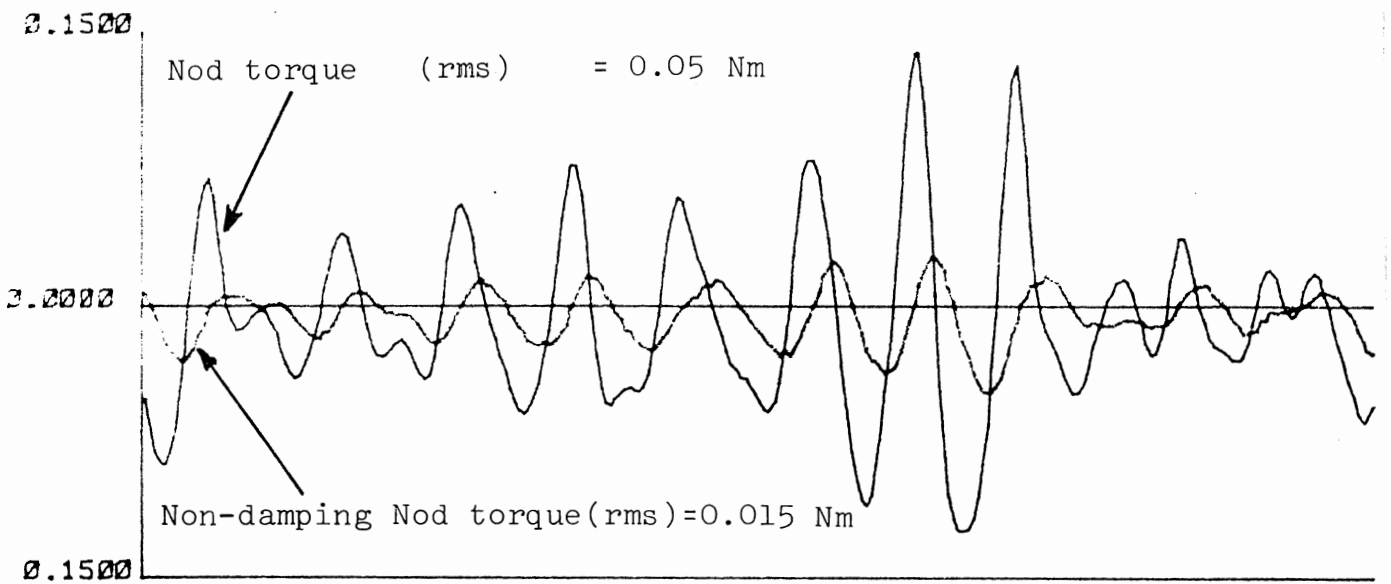
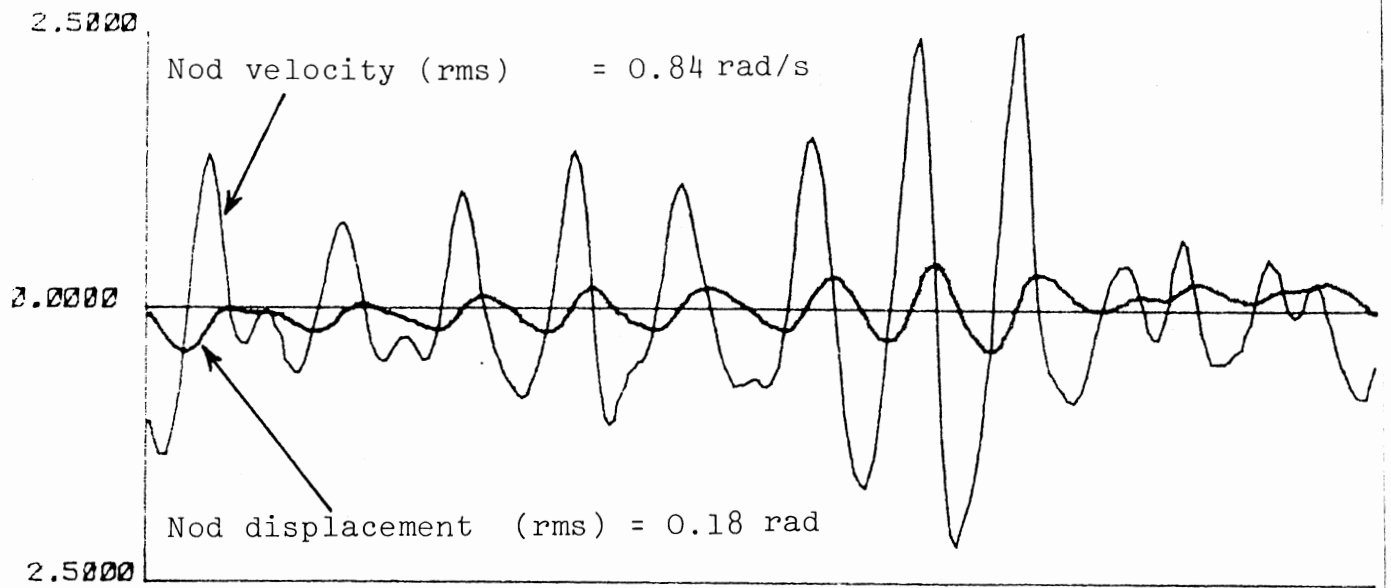Non-damping Nod torque(rms)=0.015 Nm
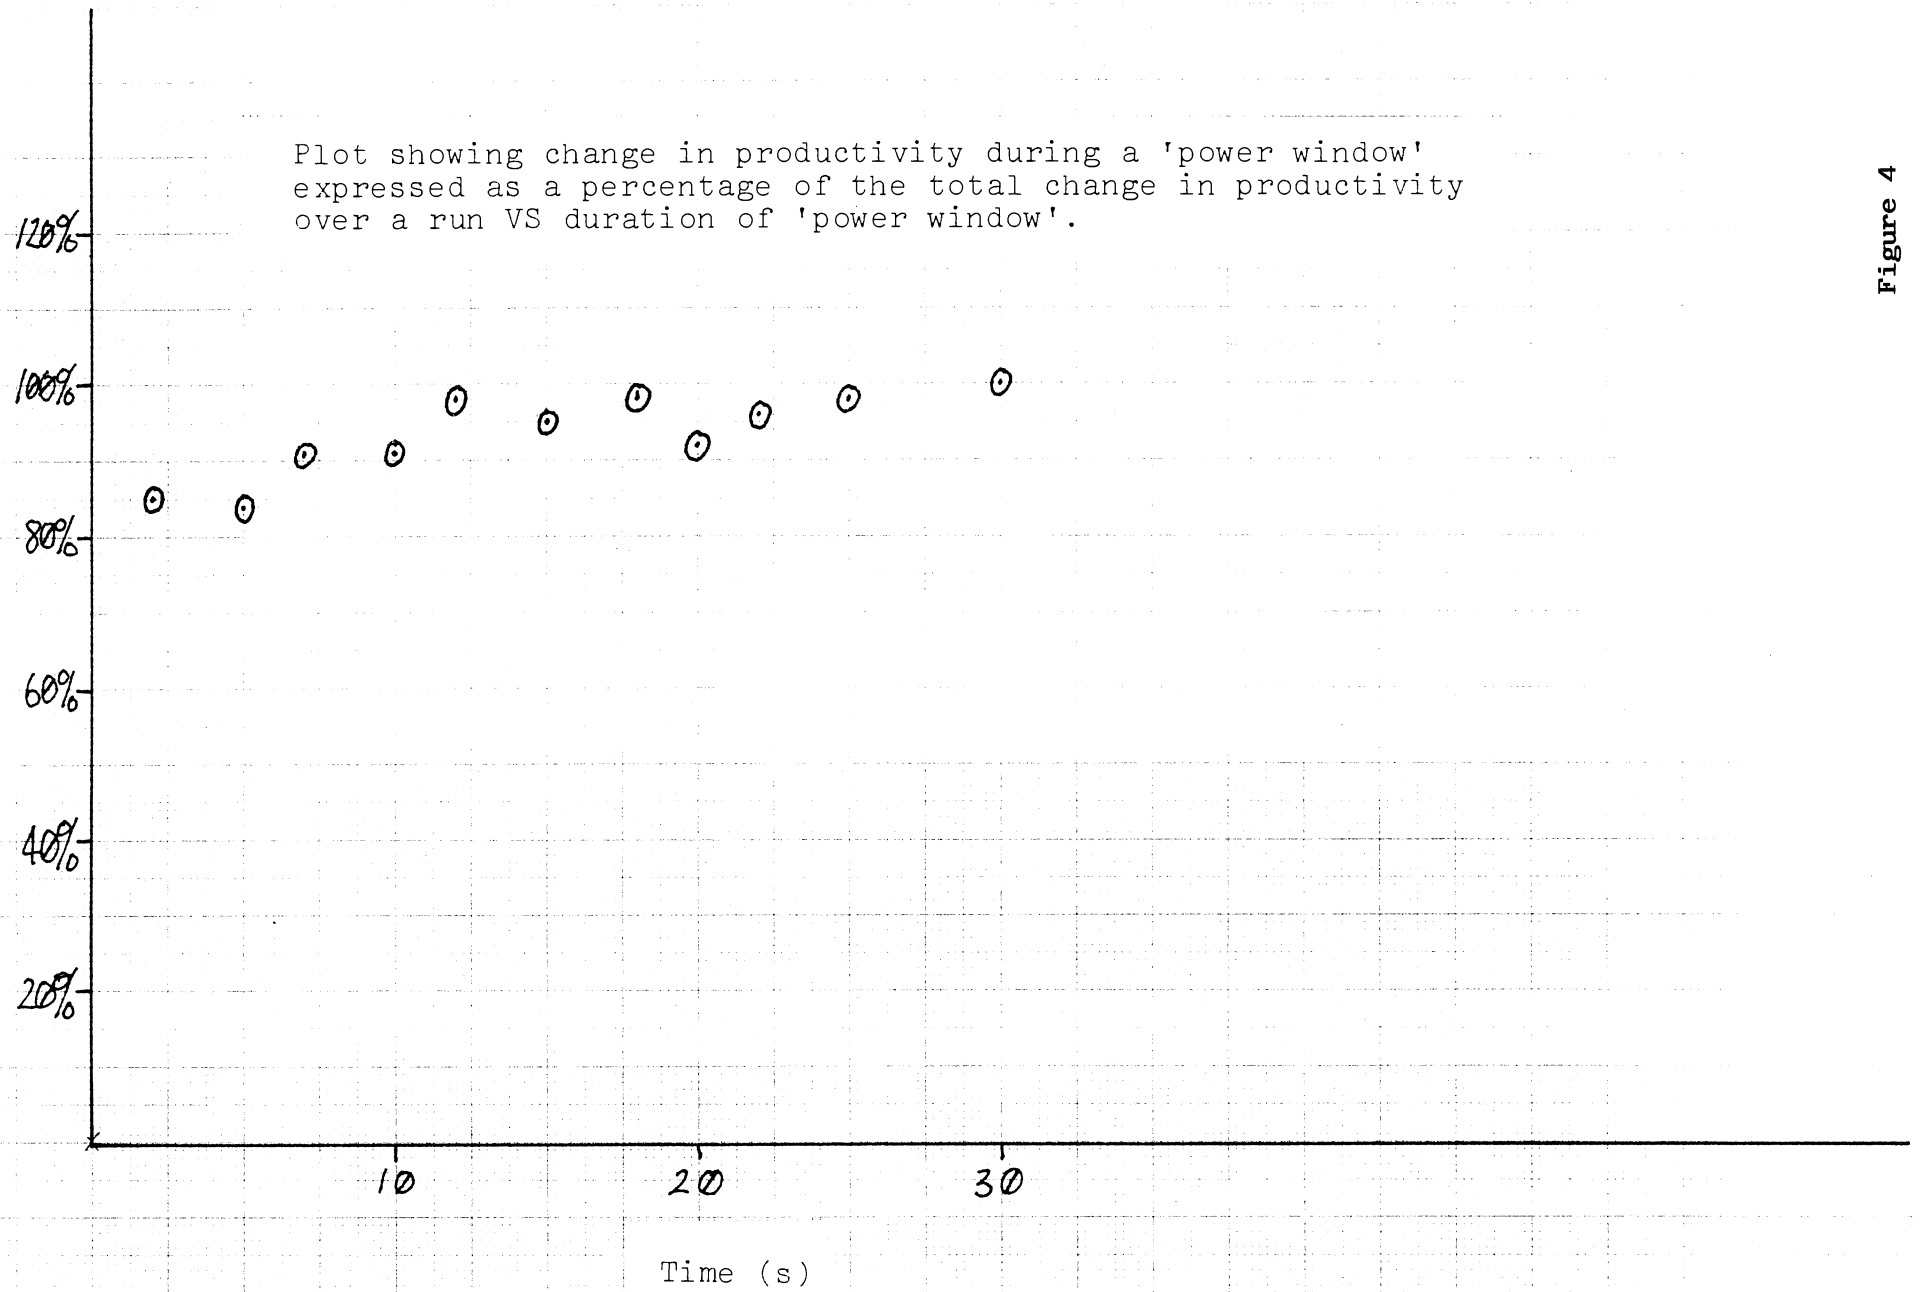
Correlation between Nod velocity and Nod torque	= 0.95
Correction coefficient	= 0.056
Original Damping coefficient	= 0.055 Nm/rad/s
Correlation between Nod displacement and non-damping Nod torque	= 0.78
Correction coefficient	= 0.068
Original Spring coefficient	= 0.1 Nm/rad

Figure 5