MASSACHUSETTS INSTITUTE OF TECHNOLOGY

A.I. LABORATORY

Artificial Intelligence Memo No. 330 December 17, 1974

LOGO Memo No. 17

A COMPUTATIONAL VIEW OF THE SKILL OF JUGGLING

by

Howard Austin

This research has as its basic premise the belief that physical and mental skills are highly similar, enough so in fact that computation paradigms such as the ones used in Artificial Intelligence research about predominantly mental skills can be usefully extended to include physical skills. This thesis is pursued experimentally by the categorization of "juggling bugs" via detailed video observations. A descriptive language for juggling movements is developed and a taxonomy of bugs is presented. The remainder of the paper is concerned with an empirical determination of the characteristics of an ultimate theory of juggling movements. The data presented is relevant to the computational issues of control structure, naming, addressing and subprocedurization.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the Laboratory's education research is provided in part by the National Science Foundation under Grant EC-40708X.

Table of Contents

	Page
1) Introduction	1
a) Thesis: Physical and Mental Skills Are Very Similar in Nature	
b) Thesis: A Computational Paradigm Can Be Validly Applied	
2) Procedural Paradigm Explained	2
a) Terminology	
b) Taxonomy definition of a theory	
3) The Experimental Setup	3
a) Role of the Teacher	
b) Use of Video	
4) The Role of Model for Physical Skills	4
a) Cascade vs. Showers: Which Model Do You Have	
b) Implications for a Theory	
5) Performance Data	8
a) The EXCHANGE Step	
b) Performance Graph	
6) Most Common Bugs	13
a) FORWARD BUG and Fix	
b) EARLY-LATE BUG and Fix	
7) Taxonomy of Juggling Bugs	17
a) Importance of A Good Descriptive Language	
b) Bugs and fixes	
8) Timing Analysis	26
a) Human Constraints	
b) Ballistic Constraints	
9) Basic Cycle	28
a) Overall Movement Pattern	
b) Reaction Theory	

10) Anticipation Theory 29

a) Frame by Frame Timing Analysis
b) Latency Considerations
c) The Cycle From the Control Point of View

11) Bugs Come in Bunches 34

a) Observation
b) Implications for Specifying Time Information

12) The Breakthrough Phenomenon 35

a) Observation
b) Implications for Naming/Addressing
c) Graphs for S1, S2, and S5

13) Its Hard to Give Yourself Advice 40

a) Observation
b) Theoretical Implications

14) Constructing Subroutines 42

a) The Main Issue
b) Empirical Data from the Juggling Case

15) Summary and Conclusions 45

16) Appendix: Protocol for S1 47

17) References 56

Acknowledgment

I would like to gratefully acknowledge the help given by Dorothy Hummel in the preparation of this paper. Without her patient assistance during many long hours of video analysis this experiment might never have been finished. I would also like to thank Karen Prendergast for her editorial comments and assistance with the majority of the illustrations, Craig Latham for his help on the remainder and Gene Horvitz for his patience in typing the original as well as countless revisions.

Introduction

This paper analyzes the physical skill of juggling from the point of view of the procedural paradigm of the newly developing field of Artificial Intelligence.

> The study of the mechanisms of physical or motor activity has been a classical theme of a wide variety of scientific disciplines including psychology and neurophysiology. The approaches employed in the investigation of sensori-motor phenomena have ranged from the very local i.e. single neuron studies to studies of global skill acquisition strategies. Important issues this research has helped to clarify include the nature of reflex and the precise extent to which sensory input controls motor activity.

Despite vast amounts of sensorimotor data however, we are still largely in the dark on important general issues, such as the difference between "physical" and "intellectual" activity as well as on specific questions such as how new physical skills are constructed from existing ones. The present research attempts to sharpen these issues by means of extremely detailed empirical observations and by the development of theoretical models. A major conclusion of the paper is that physical skills are considerably more "intellectual" in nature than hithertofore believed

The Procedural Paradigm

This section is a brief aside about the Artificial Intelligence Philosophy which underlies the experimental approach. A more extensive discussion can be found in Minsky, Papert and Winograd (1-3).

The most important aspect of any theory or paradigm is the degree to which it explains, or has the potential to explain, observed phenomena. Of course different paradigms use different criteria for testing to see whether or not a given theory "explains" a particular set of observations.

> The procedural paradigm of Artificial Intelligence (A.I.) holds that knowledge is stored internally in a form which can be modeled accurately by means of computer procedures. Hence the relevant questions for a given theory have to do with the control structure associated with the procedures, the mechanisms by which that control structure retrieves and activates procedures, the bugs (i.e. mistakes) encountered during activations and the process by which those mistakes are debugged (corrected).

Since recent work in A.I. provides a wide variety of control and activation strategies from which to choose, the most crucial questions for judging the merits of a specific theory, for example a theory of juggling, are what kinds of bugs occur when a person learns to juggle and how are these bugs fixed. In particular ,if it is possible to exhibit a taxonomy of "juggling bugs" and to describe precisely how these bugs are removed then you have given a complete theory of juggling, i.e.

TAXONOMY + FIXES = THEORY

Section 7 gives such a taxonomy for 3 ball cascade juggling, obtained by extensive analysis of video protocols of adult subjects learning to juggle. As we shall see later the precise description of a process by which observed bugs can be corrected is considerably more difficult to obtain. Before proceeding further however the details of the experimental setup are given.

The Experimental Setup

This sections explains two unusual aspects of the experimental situation. The first aspect is the extensive use of instruction, i.e. a teacher, during the learning period. The second is the use of video equipment to record the experimental trials.

The use of a teacher was very nearly mandatory given the decision to study the skill of juggling. Juggling was felt to be an excellent choice for detailed analysis since it appears to be a very complicated skill but in fact requires little more than the ability to toss and catch an object and the ability to visually track moving objects. These prerequisite skills are usually highly developed in normal adults. Hence the learning required for juggling focuses on the recombination of existing well-developed movements rather than the formation of entirely new ones.

Unfortunately, informal experiments show that very few people learn even the most simple kinds of juggling without the benefit of instruction. Although it is perhaps reasonable to conjecture that most people could eventually learn to juggle on their own <u>given sufficient motivation</u>, the amount of time required would most likely rule out short term learning experiments as well as unduly complicate the data collection problem. These problems were avoided by the active use of a teacher during the learning trials.

The use of the video equipment to record data was likewise essentially forced. It is virtually impossible to accurately record experimental trials which involve the study of complicated human movement without some sort of visual recording device such as a film or video technology. Video tape is more desirable than film since it requires no development time and can be erased and reused.

CASCADE vs. SHOWERS: WHICH MODEL DO YOU HAVE?

The following paragraphs explain in detail the precise nature of the experimental task. In addition, the first important empirical observation is presented.

The subjects were randomly selected undergraduate, graduate and faculty members of the M.I.T. community ranging in age from 18 to 50. Twenty subjects were involved of which approximately half were video taped. The rest of the experiments were recorded by a combination of third party observers and/or audio recordings. No significant difference was found in the two groups.

The experiment proceeded as follows: after a brief juggling demonstration, the subject was asked to demonstrate, in slow motion, his or her model of juggling.

> Juggling can be defined as keeping more objects in the air than you have hands. For example keeping 2 objects in the air with one hand like this

Figure 1

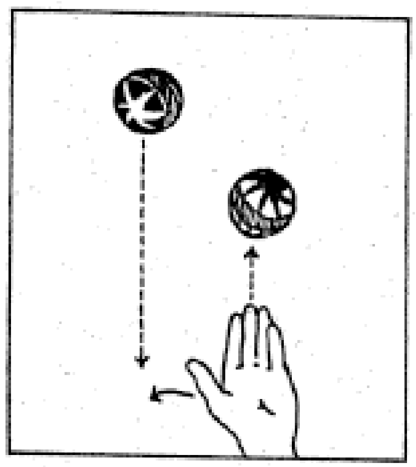

> is juggling since only one hand is involved but 2 people tossing 4 balls back and forth simultaneously are not juggling since a total of 4 hands are being used.

Cascade juggling, perhaps the most common and certainly the most simple form of 3-ball juggling is illustrated by the following sequence of sketches:

Figure 2

CASCADE JUGGLING

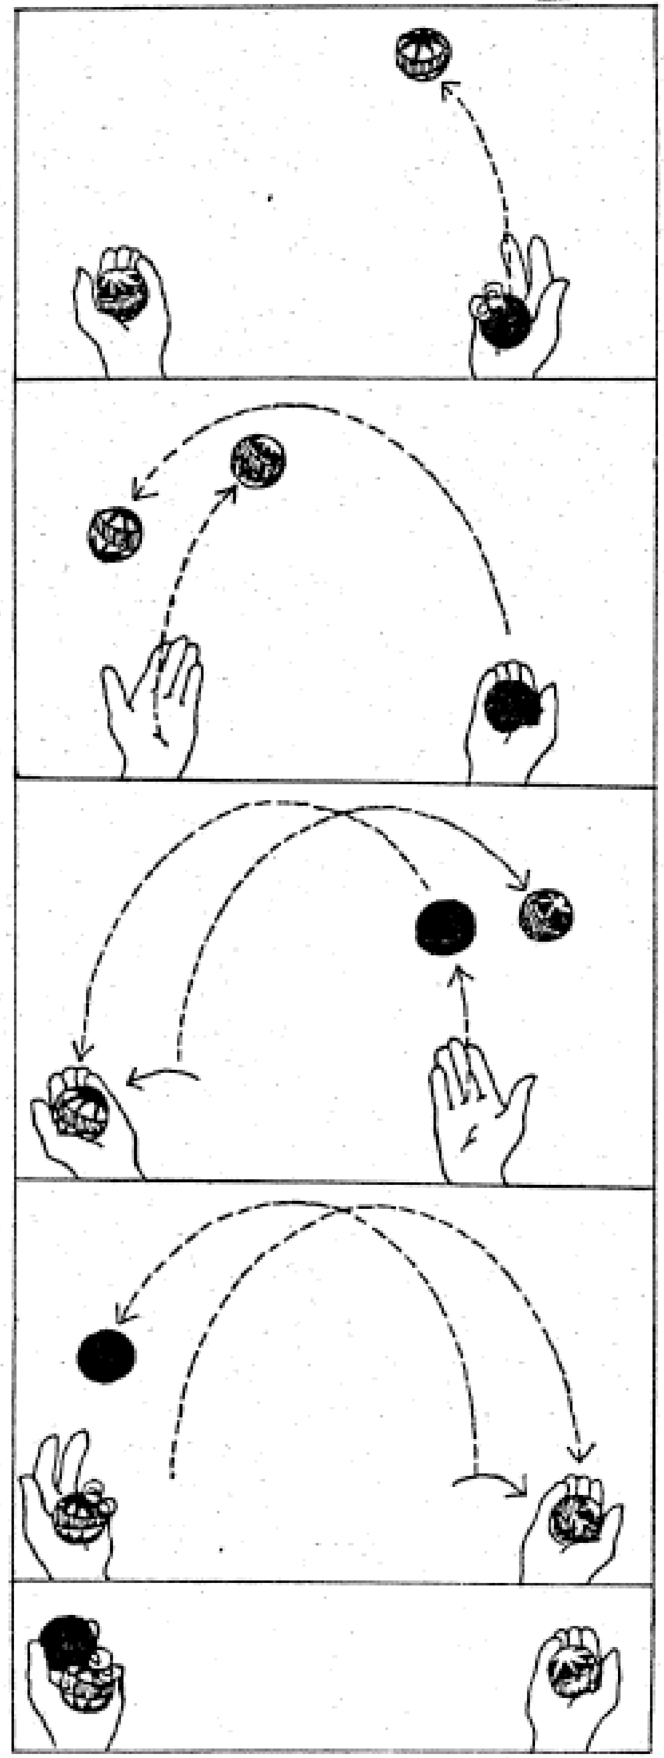

It is interesting to note that almost all of the subjects tested did not describe the "cascade" model of juggling. Rather, they almost invariably described a form of what is commonly known as "showers" juggling in which one hand does the same throw over and over while the other hand does all of the catching.

SHOWERS vs. CASCADE

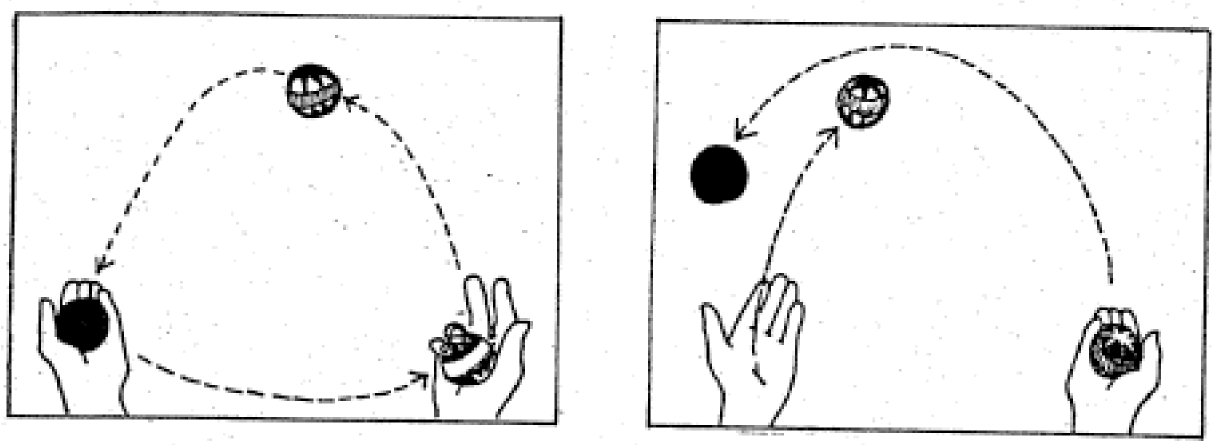

Figure 3

The difference in the two techniques can be said (at the risk of over-simplification) to lie in the fact that during SHOWERS a given hand either tosses or catches, but never both, while in CASCADE a given hand alternately tosses and catches. In SHOWERS the catching hand either physically passes each ball to the tossing hand or tosses it horizontally to the tossing hand. Whereas in CASCADE the tosses are all vertical no matter which hand they come from.*

This observation, interesting in its own right, is important experimentally because the initial demonstration each subject was given was a

*for an excellent description and further details, see CARLO (4) demonstration of CASCADE juggling which is considerably easier to learn than showers. The subjects persisted in the SHOWERS model even after the CASCADE demonstration and, left to their own devices, would have undoubtedly failed even more miserably than subjects left on their own with the CASCADE task. The important observation her is that the formation of juggling models, be it CASCADE or SHOWERS or whatever, is an entirely mental process. It has nothing whatsoever to do with muscles, neural pathways in the usual sense or sensory feedback. Hence the first step in the experiment provides us with our first clue as to the relationship between "physical" and "mental" acts. You cannot perform a physical act (at least a skilled one) until you have an accurate mental model or representation of that act!

Performance Data

This section redefines the experimental task and gives performance data for the "typical juggler" in terms of that definition.

After correcting "model errors" the subjects were led through a series of exercises in which they learned to control first one toss, then two, then three and so on. CASCADE juggling has the nice feature that both a SINGLE TOSS,

Figure 4 SINGLE TOSS

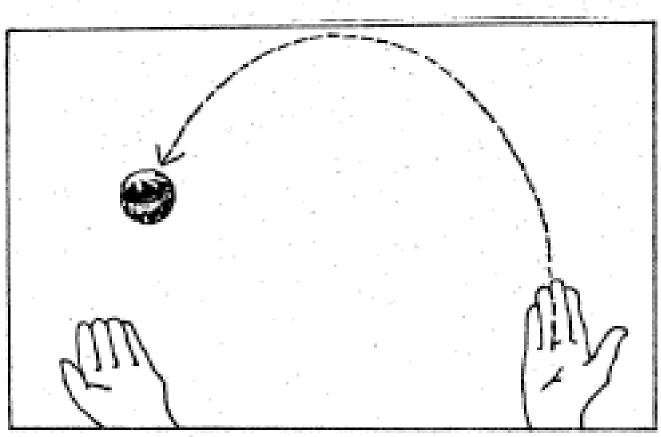

and a DOUBLE TOSS,

Figure 5

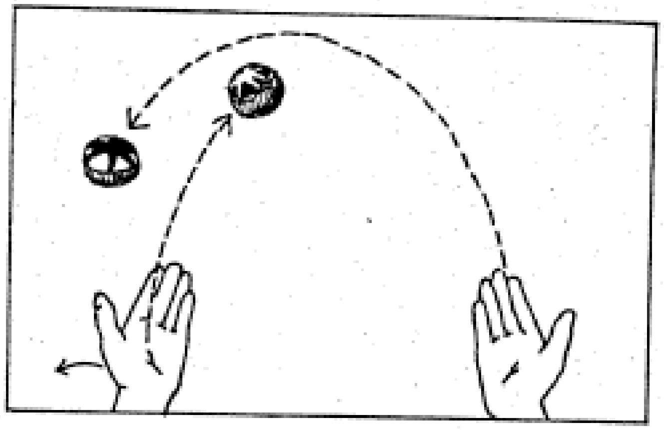

DOUBLE TOSS

are very easy to do. Almost everyone does them reasonably well on the first try. The TRIPLE TOSS is considerably more difficult and hence deserves further analysis.

The essential feature of the TRIPLE-TOSS is that it requires, after an initial TOSS with the right hand*, a rapid TOSS-CATCH sequence with first the left and then the right hand followed by a concluding left-hand CATCH. The TRIPLE-TOSS is illustrated in figure 6.

*Virtually all juggling algorithms are handwise reversable.

TRIPLE TOSS

Figure 6

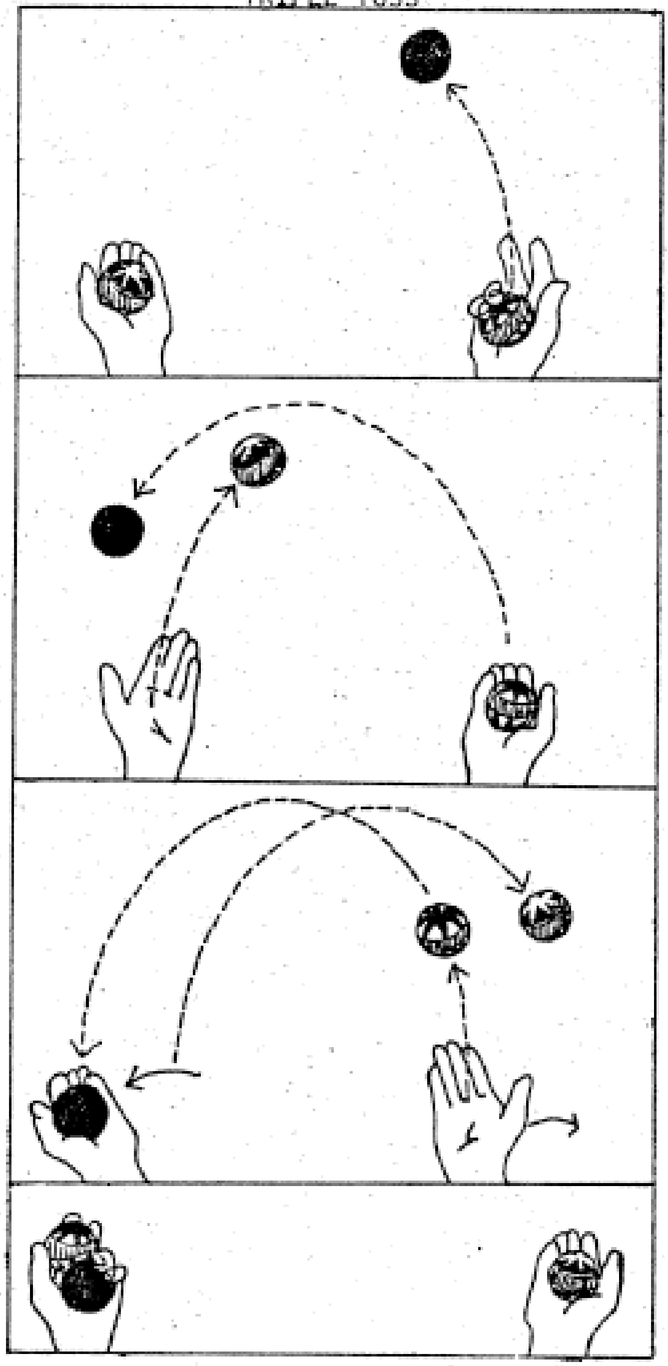

TOSS by right hand

TOSS-CATCH by left hand

TOSS-CATCH by right hand

Concluding CATCH by left hand.

The real difficulty lies in the rapid TOSS-CATCH sequences, for good reason as we shall see later. For now we shall be content to observe that if we relabel these sequences, calling them EXCHANGES, then the description of CASCADE juggling reduces to:

TOSS	to initialize
EXCHANGE	
EXCHANGE	
EXCHANGE	
.	an indefinite series of exchanges
.	
.	
EXCHANGE	
EXCHANGE	
CATCH	and a concluding catch

Given this description, it is easy to see why the ability to do 3 tosses is an important step in the learning experience. At that point the subject performs the hardest step, the EXCHANGE, once with each hand:

TOSS	
EXCHANGE	(left hand)
EXCHANGE	(right hand)
CATCH	

The rest of the learning process consists of fine-tuning the basic EXCHANGE sequence. The experimental data shows that the average person takes approximately 7 minutes to succesfully complete 3 tosses. Figure 7 gives performance data for a hypothetical average juggler in terms of number of tosses achieved after a given number of minutes practice.

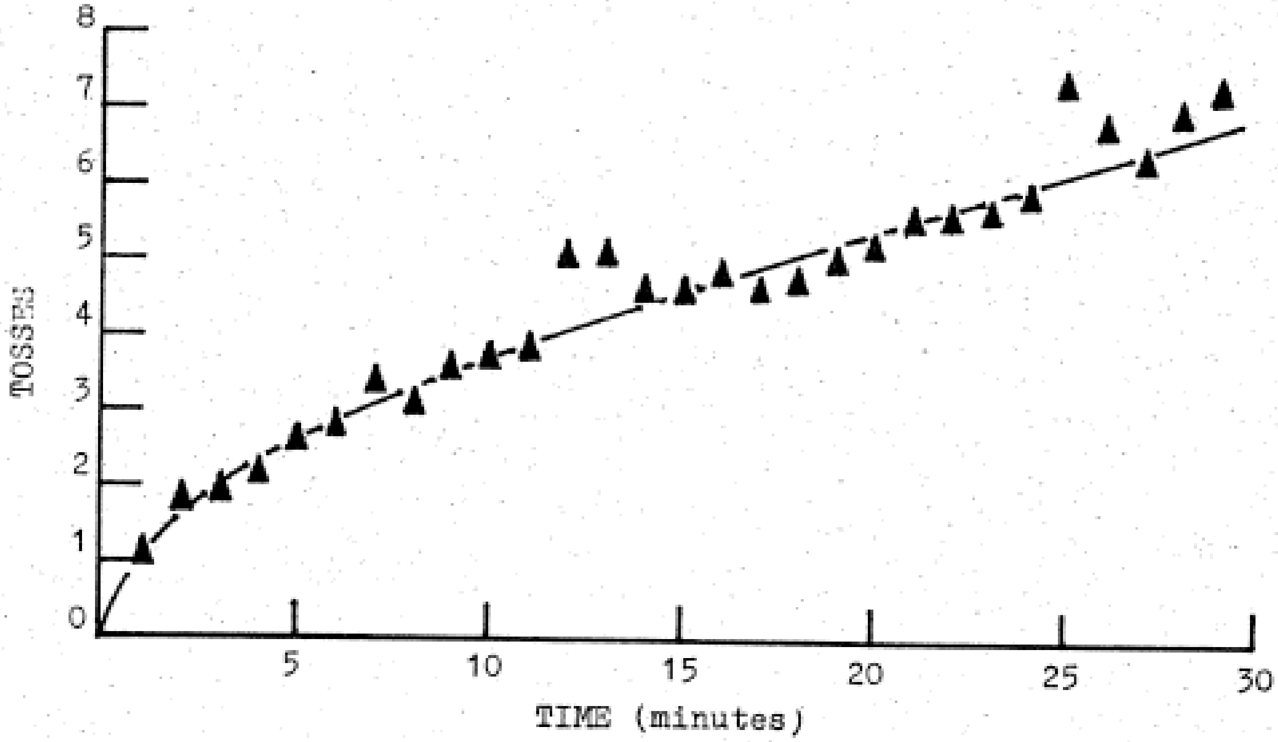

Fig. 7. Performance of the average juggler. This graph is intended to indicate the performance of a hypothetical "average juggler" during the first 30 minutes of the experimental training period. From this perspective it makes sense to say, for example, that the average juggler can do 2.5 tosses after 5 minutes. It should be noted that several of the subjects quit before the 30 minutes elapsed. This was due to (a) frustration (4 people) or (b) the achievement of 10 or so tosses (5 people). Nevertheless, it is felt that the graph is representative of average performance.

Common Bugs and Their Fixes

Most people exhibit a fairly standard sequence of bugs. The following section discusses in detail two of the most common bugs and their recommended fixes.

TIMING BUGS

If we suspend now the narration of the experiment and return for a moment to the procedural paradigm, we see that the EXCHANGE task is important for another set of reasons. It is here, at the EXCHANGE sequence, that the first serious bugs begin to show up in the learning process. These bugs usually relate in one way or another to either subtle or gross timing errors. Thus it is fair to say, at the risk of precision, that the most common kinds of bugs in learning to juggle are timing bugs.

The most frequently occurring timing bug usually manifests itself as a variation on the following theme.

FORWARD BUG

Figure 8

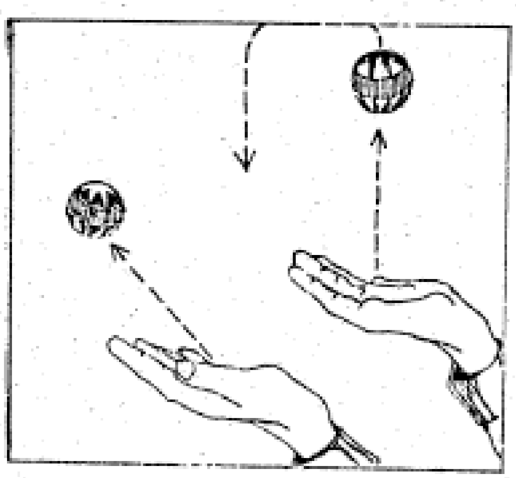

Normal right hand TOSS followed by FORWARD left hand TOSS

The first ball is tossed correctly. It has the proper height (slightly higher than the forehead) and width (slightly less than shoulder wide). The second ball, however, is tossed almost directly forward with little or no vertical velocity. This bug appears to be due to the common tendency to swing the forearm from the elbow when tossing an object. The resulting trajectory is then critically dependent on the precise instant the object is released by the fingers. An early release, perhaps caused by concern about the incoming ball, causes the second ball to go forward rather than up. A late release causes the ball to be thrown back towards the juggler.

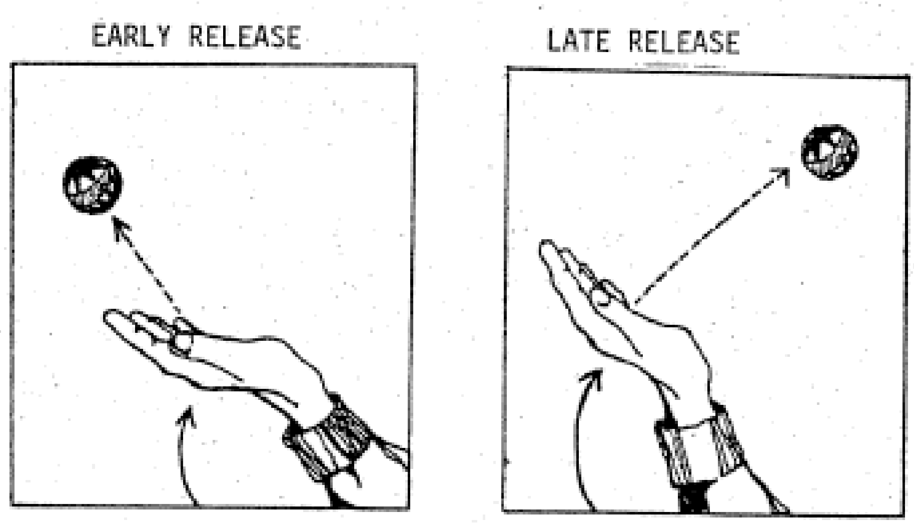

Figure 9

Of course the correct solution is either to release the outgoing ball at approximately the vertical tangent point of the swing or to move the arm in a vertical direction only so that it doesn't matter when the ball is released. The 2 "fixes" are accomplished respectively by having the subject juggle while standing near (within arms reach) and facing a wall or by teaching the subject a special largely vertical throwing motion known as "popping."

THE WALL POPPING Figure 10

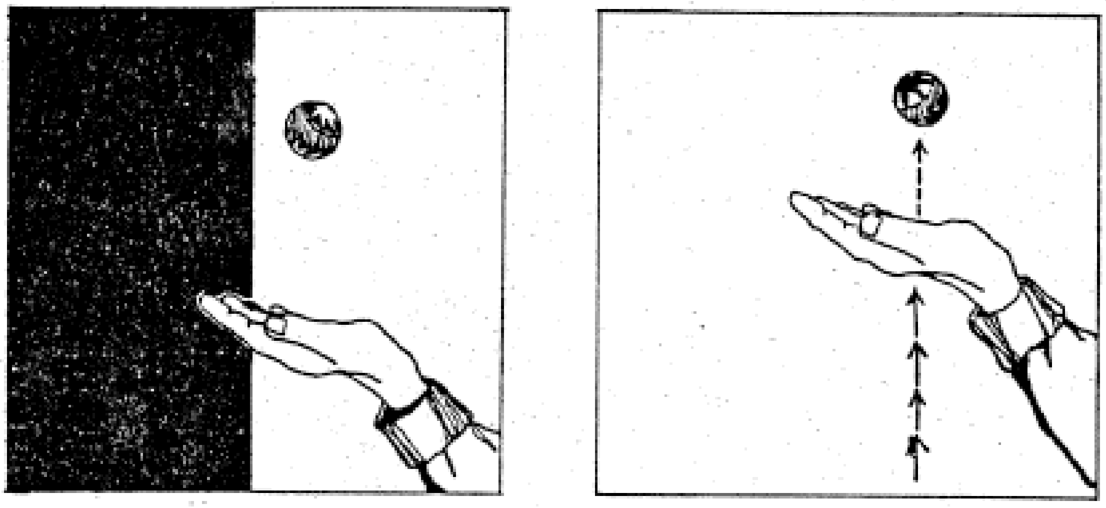

The FORWARD TOSS bug sometimes occurs on the second toss but usually appears on the third toss or at other points of difficulty.

THE EARLY LATE BUG

It should be noted at this point that the FORWARD bug is relatively ease to fix in comparison to another member of the timing class, the EARLY-LATE bug. The EARLY-LATE bug occurs as the name suggests, when one or more of the TOSSES is early or late with respect to the evenly spaced rhythm of smooth juggling. Unevenly spaced TOSSES necessitate especially rapid (or slow) EXCHANGES one step later which are usually beyond the reach of beginning jugglers.

EARLY-LATE BUG

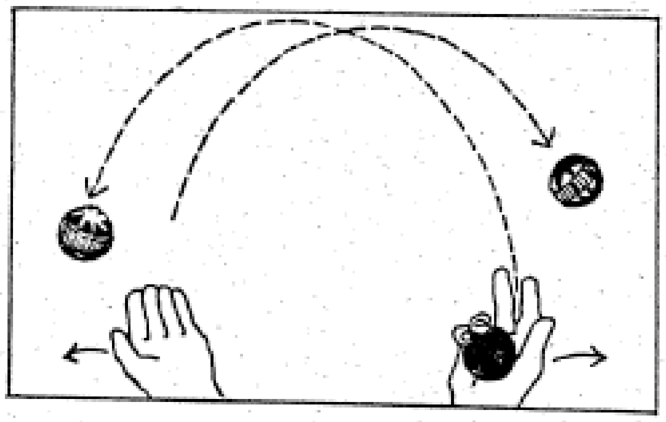

Figure 11

In this example an early second toss has necessitated a rushed right hand EXCHANGE at the same time a left hand CATCH is required.

The fixes for the EARLY-LATE bug all involve helping the subject to identify a good timing sequence for TOSSES. One way of doing this is a game called DROPCATCH. In DROPCATCH the subject is asked to execute a modified double toss in which the first ball is tossed as usual but the second ball is dropped or thrown onto the ground <u>at the last possible instant</u> before the incoming ball arrives. This of course focuses attention on the timing involved in holding on as long as possible. With this newly discovered "timing point" fresh in mind (which is very similar to the correct time for EXCHANGE initiation) it is frequently possible to achieve a good TOSS rhythm on the next trial. Perhaps a better method is to prevent the occurrence of the EARLY-LATE bug as much as possible by strongly emphasizing a proper rhythm during the demonstrations.

A Taxonomy of Juggling Bugs

This section presents the main body of empirical results produced by the experiment, namely a taxonomy of juggling bugs and their associated fixes.

The preceeding timing bugs, FORWARD-TOSS and EARLY-LATE as well as the WRONG-MENTAL-MODEL bug mentioned earlier have been analyzed in detail (along with their associated fixes) since they are representative of the whole set of bugs encountered during the experiment. This set is summarized in the taxonomy of "juggling bugs" given in figure I2. An important aspect of the taxonomy has to do with notion of a good descriptive language. In this case and in the protocol in the appendix, the names in the taxonomy attempt to be both a name and a description of the physical event. This device has two fold purpose. It greatly simplifies the difficult task of real-time analysis of complicated movements. It also facilitates communication about those movements. For more information about the important but difficult problem of developing descriptive languages for human movement see Birdwhistle (5).

Figure 12A

A TAXONOMY OF JUGGLING BUGS

TOSSING BUGS

NAME	EXAMPLE	FIX
FORWARD	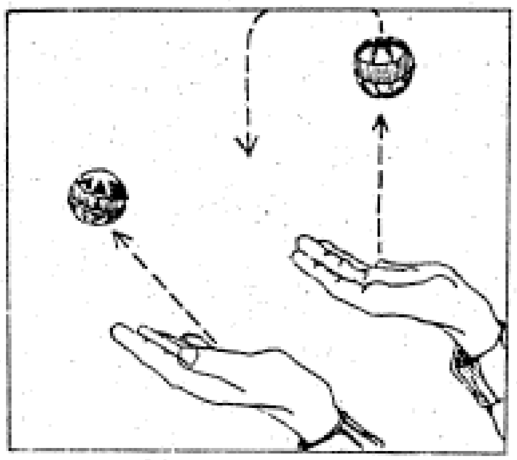 VERY FORWARD TOSS	a) FACE A WALL b) POPPING c) DROP CATCH
BACK	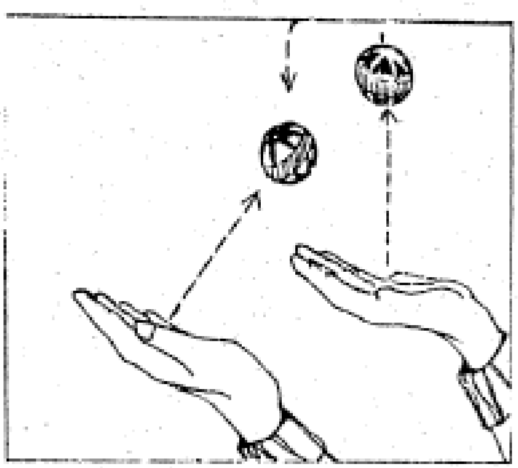 SLIGHTLY BACK TOSS	POPPING (when it occurs on the first toss as is most common). Later occurrences are usually caused by a gross error on preceeding cycles and hence disappear when the prior error is fixed

Figure 12B

WIDE

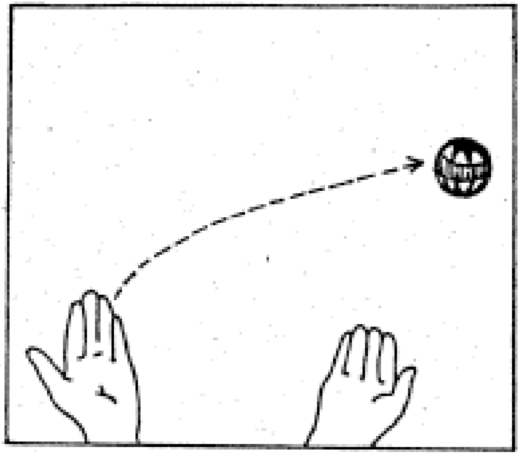

NARROW

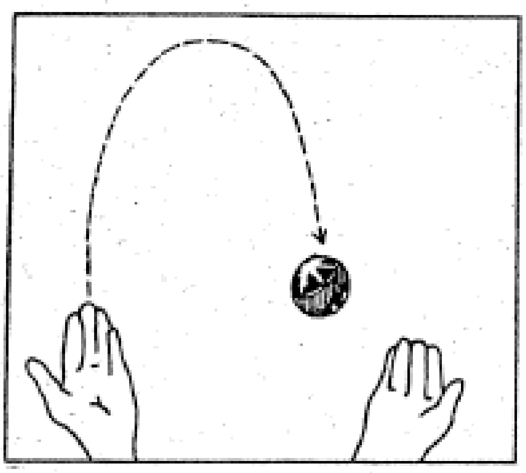

VERTICAL

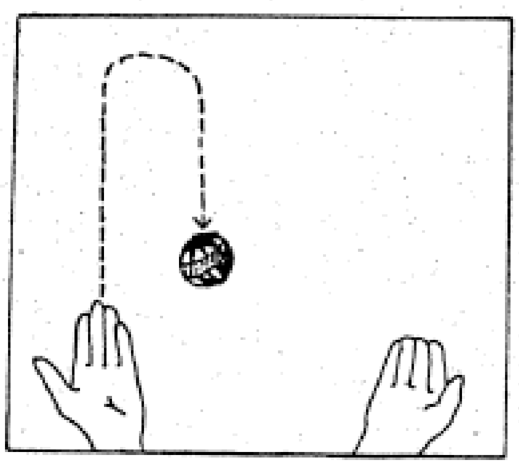

The next three bugs usually occur either as the result of a bad model of what a toss should look like or as the consequence of an earlier mistake somewhere in the preceeding cycles. In the first case the fix is a careful demonstration of a good toss along with verbal guidelines like "head high" or "shoulder width."

In the latter case they usually disappear along with the prior mistake. (The VERTICAL bug has a separate name because it is so disruptive when it occurs that it deserves special attention).

Figure 12C

COLLISION

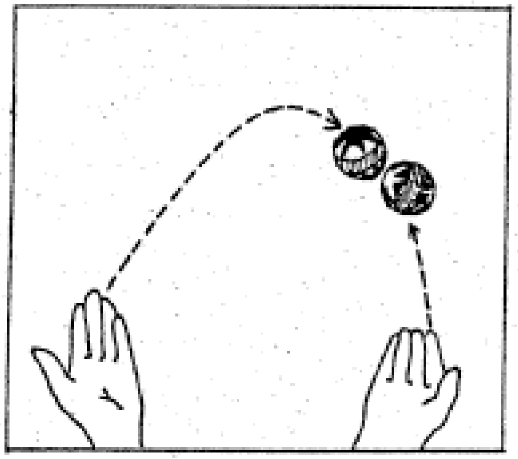

a) Demonstrate the difference between underthrowing and overthrowing.
b) Pretend there's a large pipe beside and "inside" the trajectory of the incoming ball into which you must throw the outgoing ball.

CAN'T RELEASE

This bug most frequently occurs on the 4th toss and may perhaps be due to the pedagogical sequence used which emphasizes "getting to 3" or to an overly cautious "dynamic lookahead" ability which predicts disaster on the upcoming toss. In any event the fix is to force the occurrence of the inhibited toss by either changing the toss cue via instructions like "don't throw until you hear me shout NOW" (this is a variant of DROPCATCH) or by removing the fear of failure by instructions like "don't worry about catching it for now, just throw it."

CATCHING BUGS

Figure 12D

NAME	EXAMPLE	FIX

FINGERNAIL

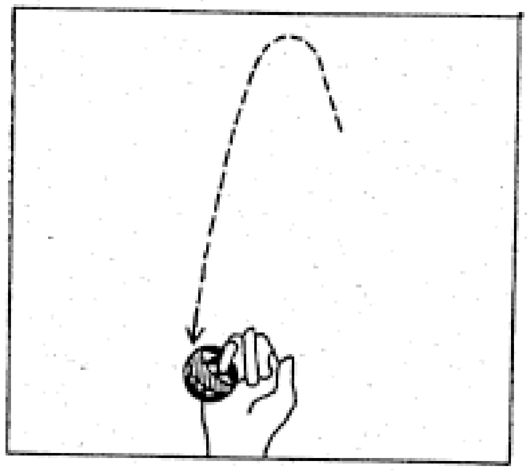

The catching errors shown here almost always stem from poorly executed previous toss or sequence of tosses. Hence the fix in most cases involved correcting the offending TOSS.

HANDSLIP

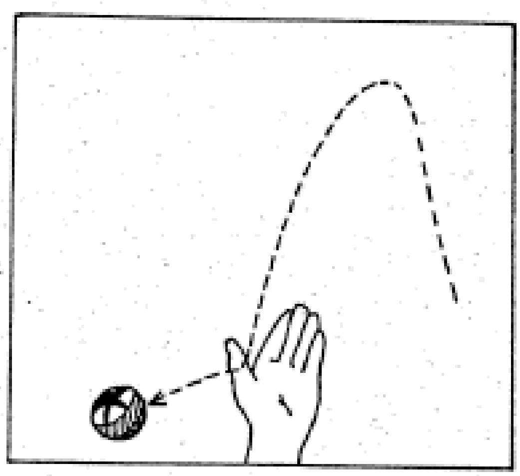

Figure 12E

CHESTCATCH

(Usually found only after serious BACK or timing bugs)

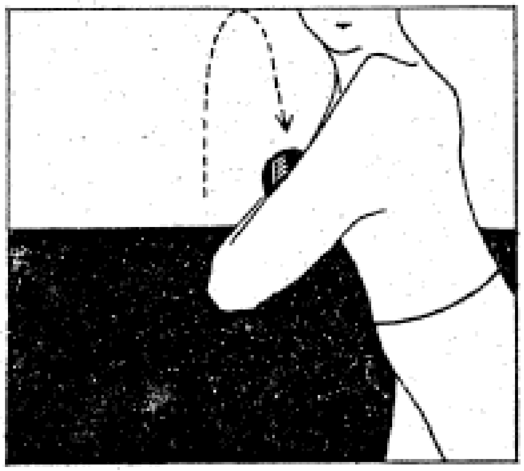

SWIVELCATCH

(severe hip rotation but no step)

LUNGECATCH

(involved one or more steps, lunges, etc. i.e., a last ditch attempt)

TIMING BUGS

Figure 12F

NAME	EXAMPLE	FIX
EARLY	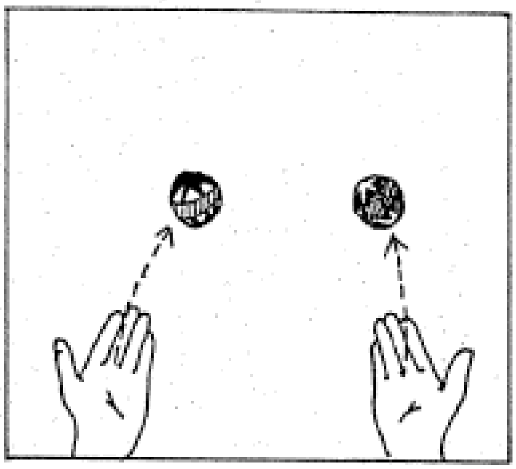	Although all of the bugs exhibited so far are, at some level, accurately described as timing bugs, some are more directly concerned with the sequencing of major events, such as tosses, rather than details within a given event such as a release within a given toss. Hence they deserve special attention and names. The fixes in all cases involve establishing a good major event sequence by games like DROPCATCH, counting out loud, demonstrations, etc.
LATE	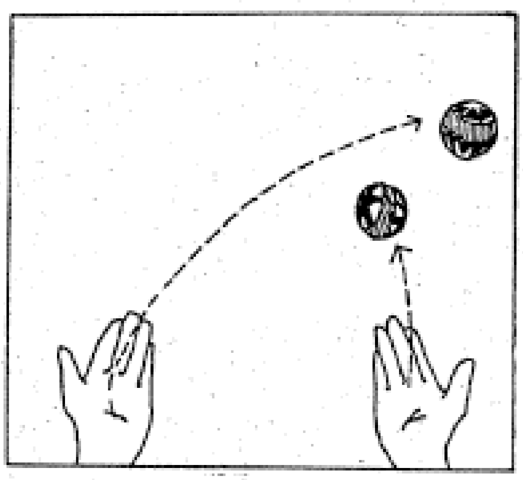	

TIMING BUGS (cont.)

Figure 12G

SLAP

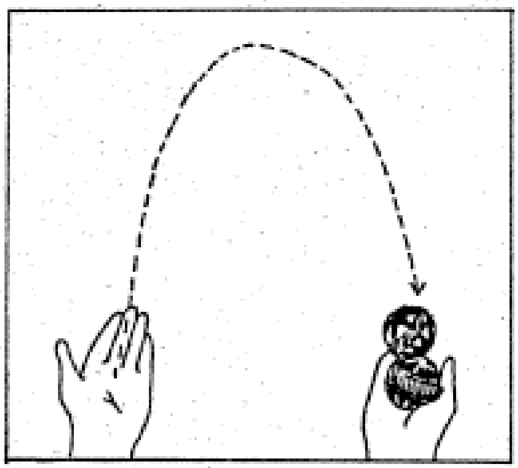

In the SLAP bug the outgoing TOSS is so late that there is no time left to throw. In a desperate attempt to both throw and catch at the same instant the hand compromises both tasks by "slapping" (while still holding onto) the supposed outgoing ball into the incoming ball. The fix for SLAPPING is to forbid the subject to do it (since it is a voluntary act). The ball must be either caught or allowed to drop untouched. This helps to focus attention on the proper EXCHANGE initiation point.

The Outlines of A Computational Theory of Juggling

Overview

The preceeding sections of this paper were concerned with the presentation of empirical data about the skill of juggling approached from a computational point of view. From that point of view the experimental observations with respect to the role of model in physical skills, the most common and most severe bugs and associated fixes, as well as the notion of a complete taxonomy of bugs all have to do with what computational experts would call software aspects of juggling. By definition however, the computational paradigm consists of a set of software instructions (by which the programmer fulfills his model of the problem, creates bugs, etc.) and a set of hardware devices by which those instructions are interpreted and executed. For people of course the hardware devices by which movement instructions are executed are muscles, neural networks, the motor cortex, and so on.

It is not possible at present to give a complete theory of this hardware device for physical "commands." It is difficult enough to give detailed theories about very small components of what might be called the "physical computer" part of the human body. The rest of this paper is concerned with establishing empirical constraints on the range of theories which might be offered for the various components of such a computer. The issues considered have to do with timing constraints, potential control structures, naming and the construction of subroutines.

Timing Analysis

It should be clear, computationally speaking, that juggling is a real time process and further, that the constraints imposed by the physics involved are fairly severe. The following discussion has as its goal a clear formulation of the timing constraints faced by an accomplished juggler as he performs his routine. The constraints derived here combined with what is known of human response times should be of considerable value in designing a control structure for a theory of juggling movements.

Human Constraints

It is apparently true that any willed skeletal movement takes approximately 200 ms. to initiate. That is to say, when your head tells any part of your body to move, including the eyes, on the order of 200 ms elapse before the movement actually begins. Although there is some variation depending on the task, the individual, and to some extent the amount of practice, the 200 ms elapse latency is essentially invariant, particularly for large movements of the kind involved in juggling. Likewise the maximum acceleration for the hand is approximately 2400 in./sec^2.

> These parameters can be verified empiricially by the following two experiments:
> Try to catch a ruler dropped (by someone else) between your almost closed fingers. The zero mark should be even with the top of the hand. The hand should be resting on some fixed object so it can't "chase" the

ruler. The fingers will usually grasp the ruler somewhere between the 7 and 8 inch marks which corresponds roughly to a 200 ms delay. This explains why it is nearly impossible to a catch a 6" dollar bill in the popular version of the experiment.

The second experiment involves trying to swat a ruler dropped past some flat surface like a table. The hand is positioned 3" from the edge of the table and moves, when the ruler is dropped, in a direction which is horizontal to the table. The ruler is held so that the end is even with the top of the table. With a little practice you can just nick the end of the ruler as it goes by. Since the ruler falls 12" in ≈250 ms, the hand moves 3" in 50 ms, assuming a 200 ms latency. This corresponds to an acceleration of 2400 in./sec.^{2}. These examples are due to Lacy (6).

<u>Ballistic Constraints</u>

Given that the hand operates with a 200 ms latency and can accelarate at no more than 2400 in./sec.^{2}, how do these constraints compare to the ones imposed by the laws of physics? The first step is to determine the velocity of the ball in normal catching regions. For a person 6 feet tall, a toss typically rises to a maximum height of 6 1/2 feet and is usually caught at slightly more than shoulder level. This situation can be modeled from the timing point of view as a ball dropped from a height of 6 1/2 feet and caught at a height of 5 1/2 feet. Thus the ball reaches a velocity of 96 in./sec. and hence requires 250 ms to travel the 12 inches from the top of its arc to the level at which it usually caught. If for some reason the distance traveled becomes as much as 24 inches the velocity of the ball increases to 136 in./sec. so that only 350 ms are required to cover the longer distance.

Basic Cycle

It should be noted that the hand which is about to catch the incoming ball (i.e. in our model the dropped ball) is currently occupied with tossing the outgoing ball and hence at the conclusion of that toss is very close to where it should be in order to catch the incoming ball. In fact if attention is focused on say the jugglers left hand the following pattern of movements can be observed.

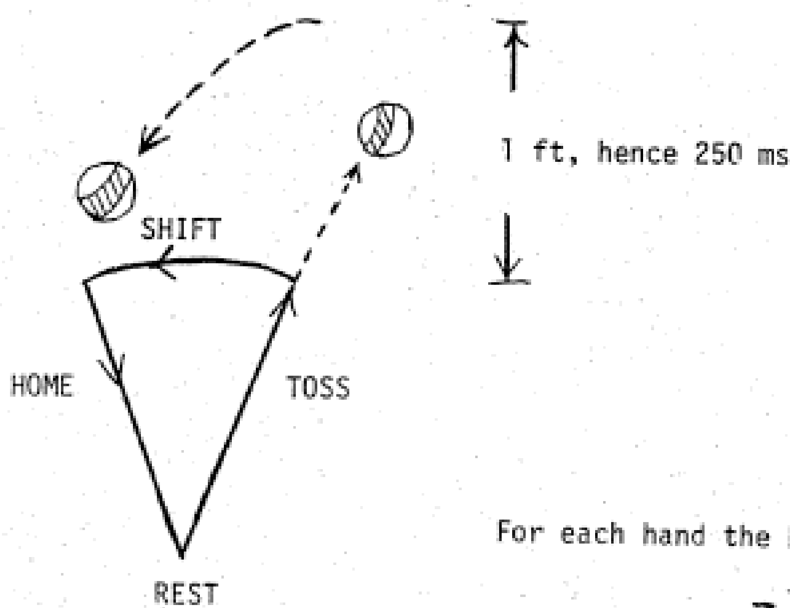

Figure 13. The Basic Cycle

For each hand the basic cycle consists of

TOSS
SHIFT
CATCH
HOME
REST

and the shift part of that cycle has to cover only 3 or 4 inches for most tosses. By the second ruler experiment it seems reasonable to believe the hand can cover such distances within the available 250 ms and hence that a kind of control structure one might call a "reaction" theory could support juggling movements satisfactorily. Such a theory would claim, for example, that at the end of the TOSS movement the eyes perceived the position of the incoming ball and "reacted" by issuing a command to move the hand under the ball. Likewise the actual contact of the hand and the ball causes another reaction in the form of a GRASP movement. The central tenent of a reaction theory of movement control is that the commands which initiate movements are issued in response to rather than in anticipation of external events.

An Anticipation Theory

Unfortunately the preceeding analysis does not include latency times for the various movements and as yet says nothing about the control of eye movements. These considerations make it obvious that a more complicated "anticipation" theory will be required for any control structure the overall theory might select. For example, figure 15 presents a frame by frame analysis of a brief (8 toss) juggling sequence. This analysis shows that the most crucial part of the EXCHANGE, the SHIFT movement, lasts on the

Figure 14

Basic Cycle
Timing

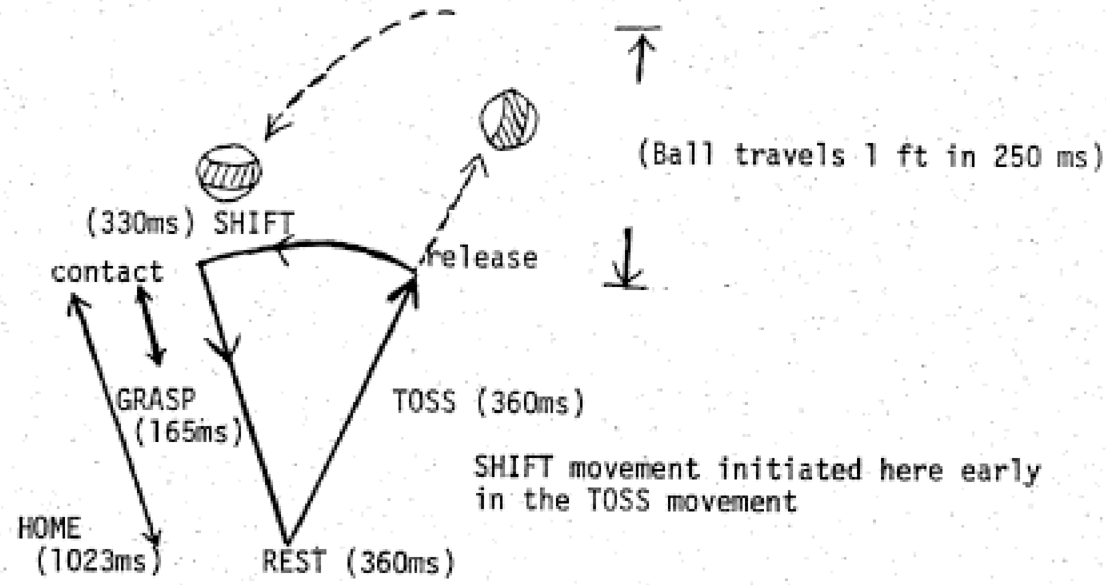

average 330 ms after it begins. Of course there is a theoretical question as to whether or not the 330 ms consists of a 200 ms wait and a 130 ms move or consists entirely of a 330 ms move. The empirical evidence from the video tape is clear on that point however. There is no observable pause at the end of the TOSS movement and certainly none of 200 ms duration. Rather the whole TOSS, RELEASE, SHIFT, CONTACT, HOME sequence appears to be one smooth movement.

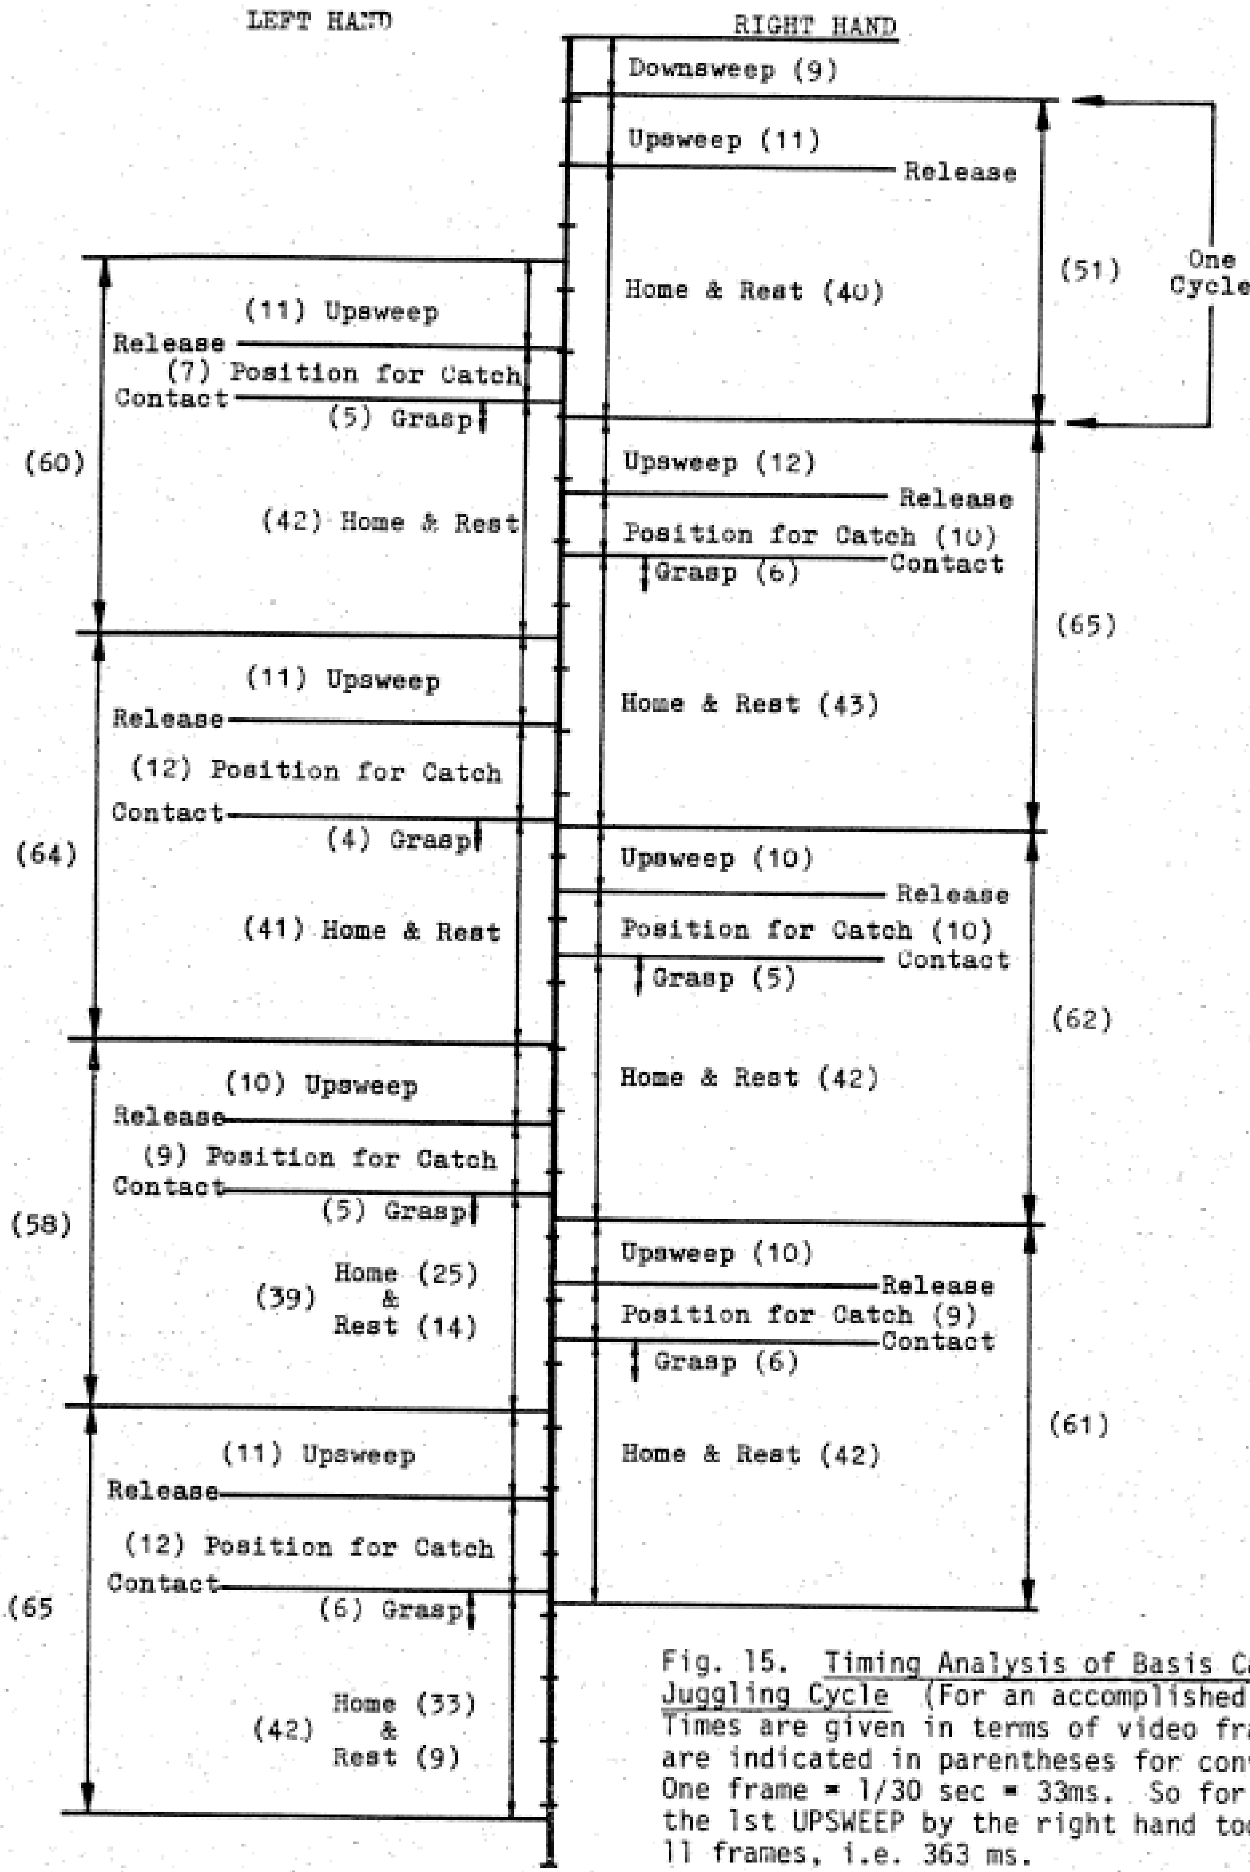

Fig. 15. Timing Analysis of Basis Cascade Juggling Cycle (For an accomplished Juggler) Times are given in terms of video frames and are indicated in parentheses for convenience One frame = 1/30 sec = 33ms. So for example the 1st UPSWEEP by the right hand tools 11 frames, i.e. 363 ms.

The only way this could have occurred from a control point of view is for the SHIFT movement to have been initiated some 160 ms into the TOSS movement, the GRASP movement to have initiated 130 ms into the SHIFT movement, etc. (see figure 16).

To return to the theoretical discussion, the issue becomes forced when eye movements are considered. At least one eye fixation is required to determine the trajectory of the ball in flight. Yet eye movements also have a 200 ms latency. So seeing the ball and causing the arm to catch it involves at least a 400 ms delay assuming an instantaneous trajectory calculation and zero time for the arm movement. But the model allows only 250 ms for the ball to fall the 12 inches involved in a normal catch. Doubling this distance adds only 100 ms or so to the time available. Thus it is clear that the SHIFT movement is anticipatory in nature in the sense that it must have been "planned," i.e., initiated, during a prior movement. Notice that the preceding discussion does not consider the possibility of quick-reaction low-level feedback loops which obviously could affect the details of the analysis. Nevertheless the primary conclusion still stands. Any theory of juggling movements must contain a control structure which is capable of supporting parallel, real-time, events of an "anticipatory response" nature.

FIGURE 16

THE BASIC CYCLE FROM THE CONTROL POINT OF VIEW

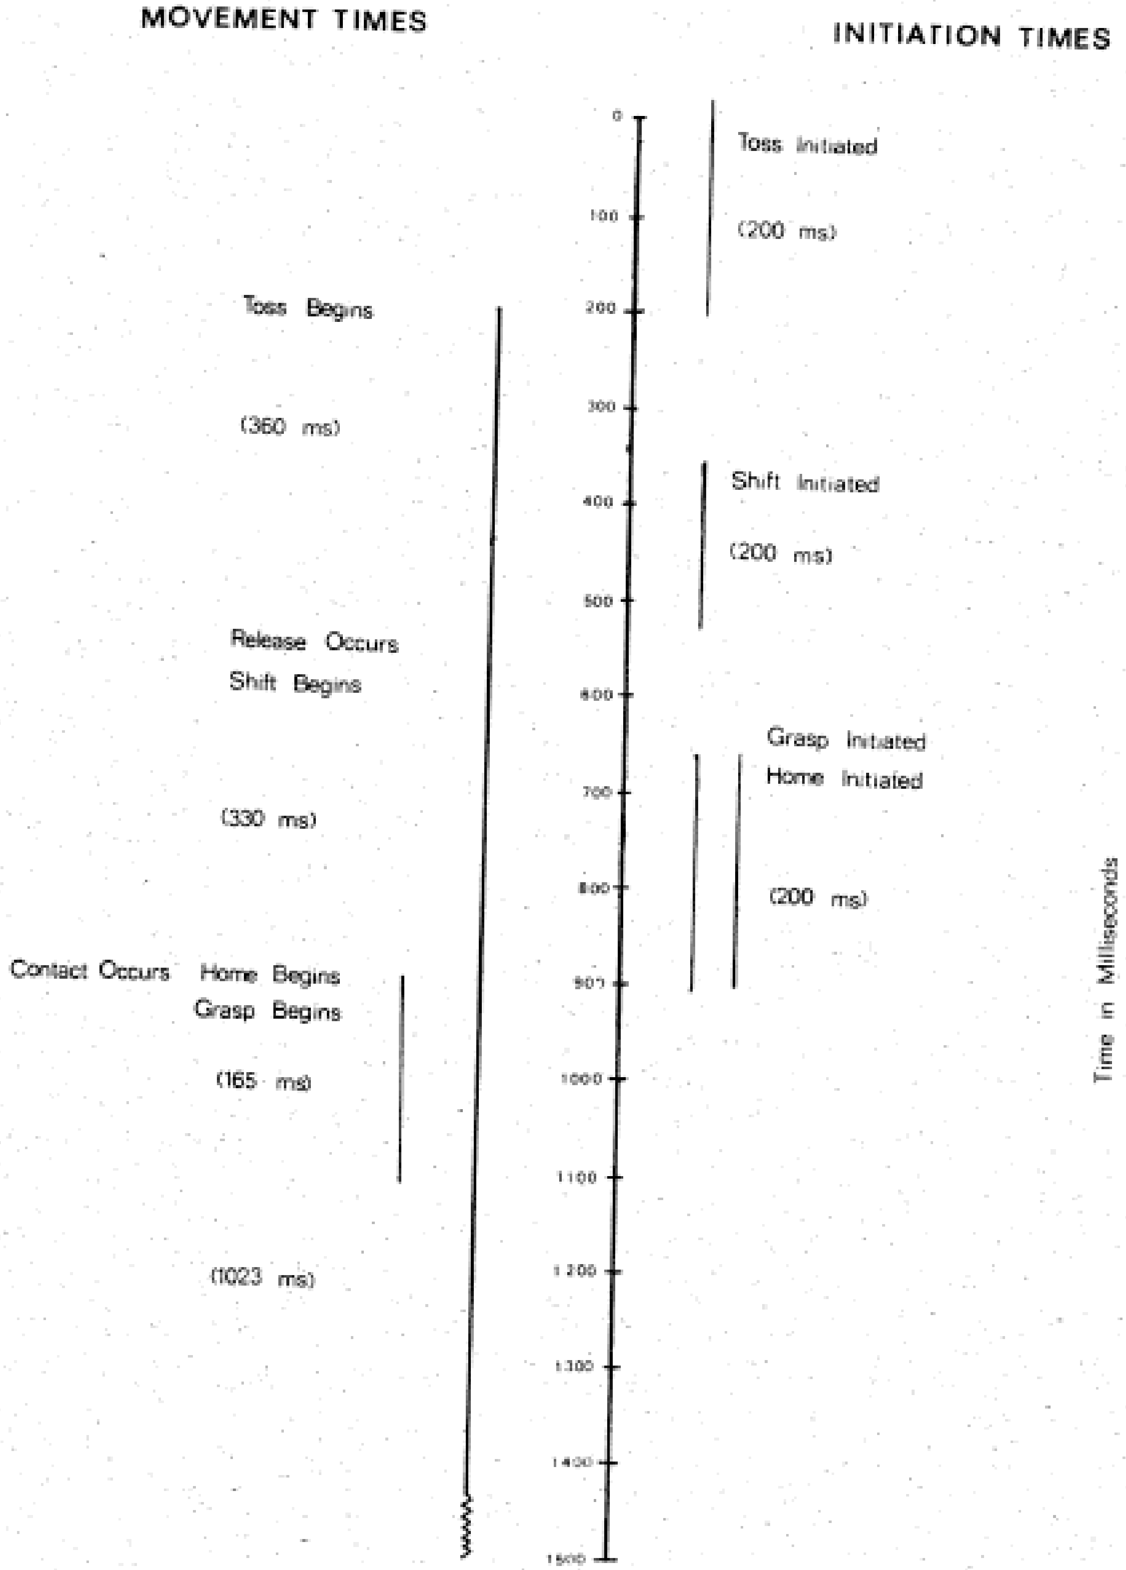

Bugs Come In Bunches

The next experimental observation also has to do with timing information but in a dffferent sense. The observation is that bugs apparently come in bunches, at least that is true for the juggling variety anyway. Its probably safe to say that above the DOUBLE TOSS level most of the subjects were initially grappling with two or more simultaneous bugs. One frequent combination was EARLY, LATE/FORWARD. The protocol in Appendix I contains many other examples.

This observation is significant for both theoretical and pedagogical reasons. One of the most important differences between mental and physical events is the fact that all physical events reduce ultimately to some pattern of muscle start and stop times. Although it could likewise be said that all mental events reduce ultimately to something like current flow,in some deep sense it seems to be true that physical events are considerable more time-dependent than mental events. We have seen for example the whole taxonomy of bugs which can be generated from subtle shifts in the timing of events in the basic juggling cycle. It is perhaps reasonable to conjecture then that simultaneous bugs are due to a combination of timing mis-specifications in the learning process and the narrow range of permissible specifications afforded by physical constraints. The important theoretical implication is that any theory of movement must have a mechanism, perhaps a whole language, which is capable of describing and modifying timing events at the millisecond level.

Pedagogically the observation is important because if, as conjectured, juggling is a reasonable representative of physical skills in general, then the indicated experimental direction is the creation of learning environments in which the same bugs occur, but one at a time. For juggling, this might take the form of a machine which allows EXCHANGE practice in slow motion, or as a separate entity. EXCHANGE bugs could then be dealt with one at a time independently of the rest of the process.

The Breakthrough Phenomenon

The single most surprising discovery made during the experiment was that for most individuals progress, as measured by either the number of the quality of successful tosses, was not at all continuous but, rather, appeared to consist of a series of breakthroughs or plateaus. These breakthroughs were not initially stable, frequently being attainable for only one or two trials, but became more frequent and stable over time. The really astonishing piece of information is that the magnitude (measured in tosses) of the breakthrough was usually 2 - 4 times that of the previous level! Jumps from 3 tosses to 10 or 11 tosses were common. See for example the sixth minute of S1 or the eleventh minute of S2 in figure 17. Almost every subject in the experiment exhibited a breakthrough which virtually doubled the previous best effort.

One explanation of this observation is that several bugs vanished or were temporarily overcome simultaneously thereby allowing a dramatic breakthrough to occur. The fact that the breakthrough is not initially stable but improves with time suggests that some form of trial and error or hillclimbing search is going on at the lower levels of timing specification. This argues against a strictly symbolic naming/addressing mechanism in the overall theory.

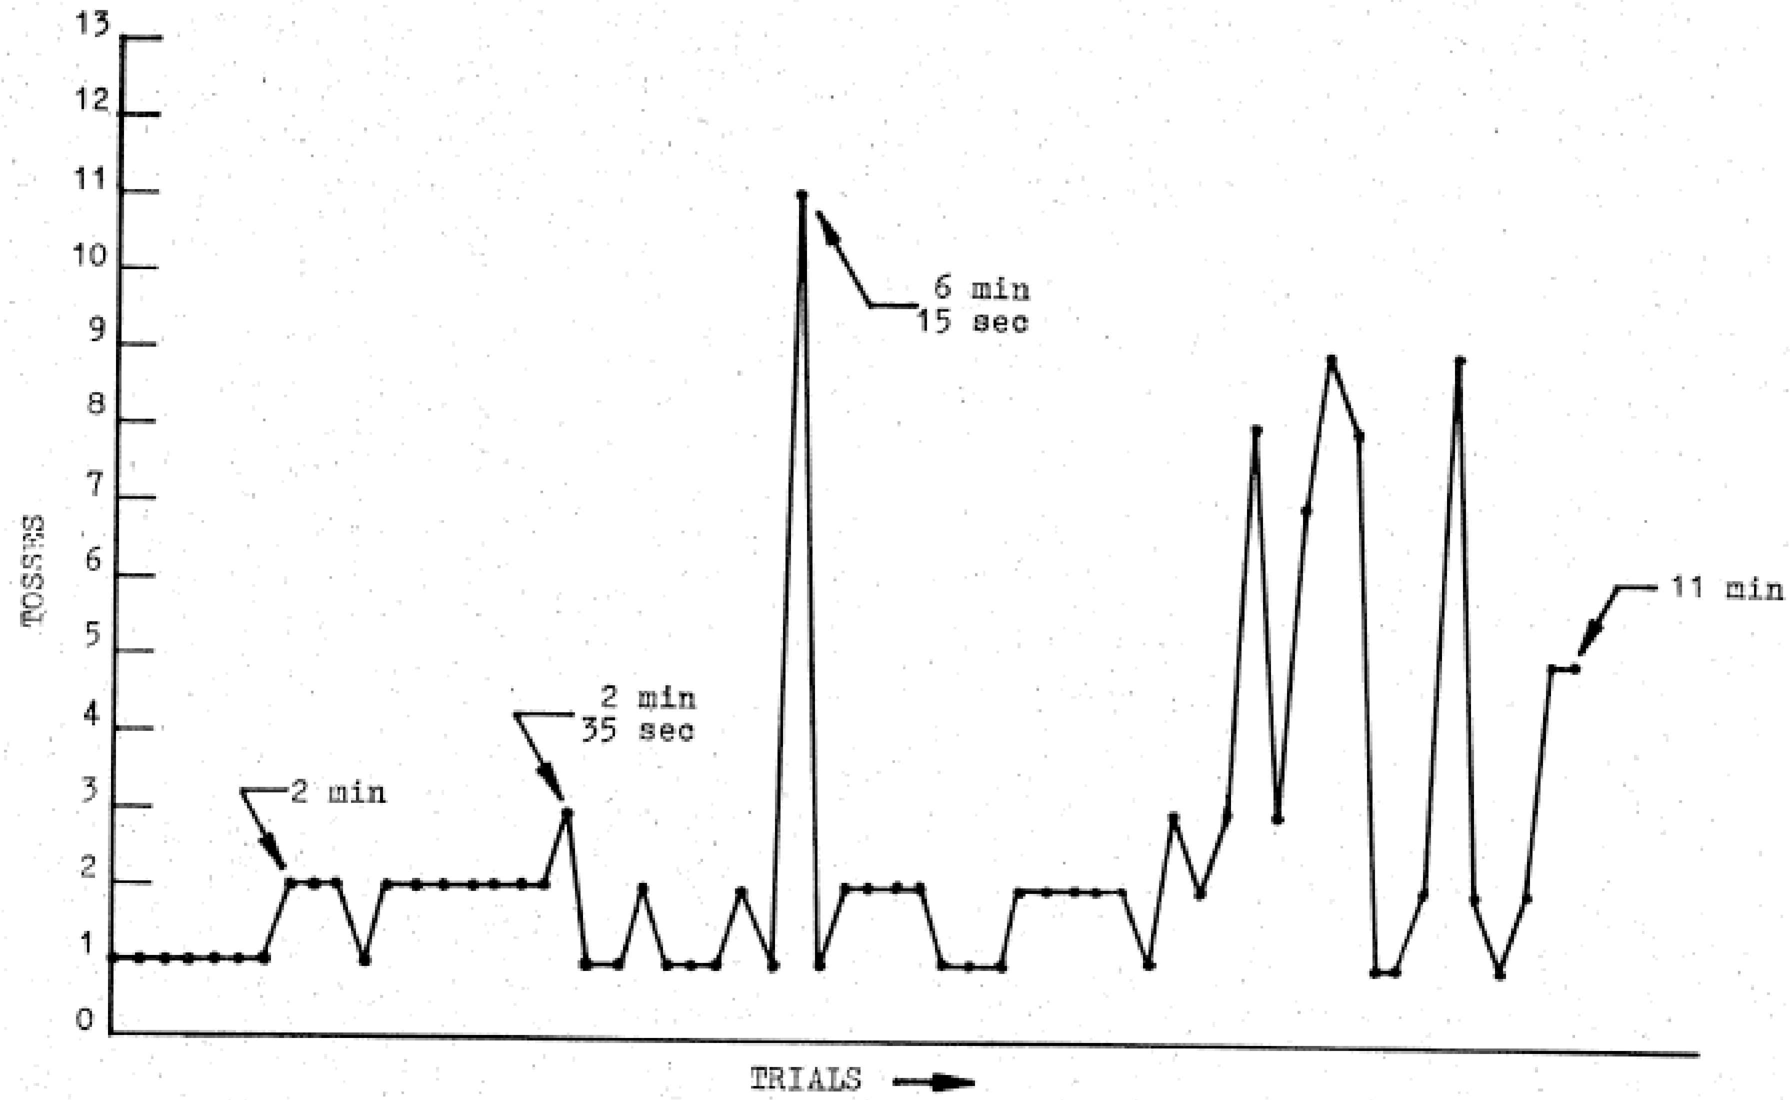

Fig. 17A. The Breakthrough Phenomenon (S1)

Fig. 17B. The Breakthrough Phenomenon (S2)

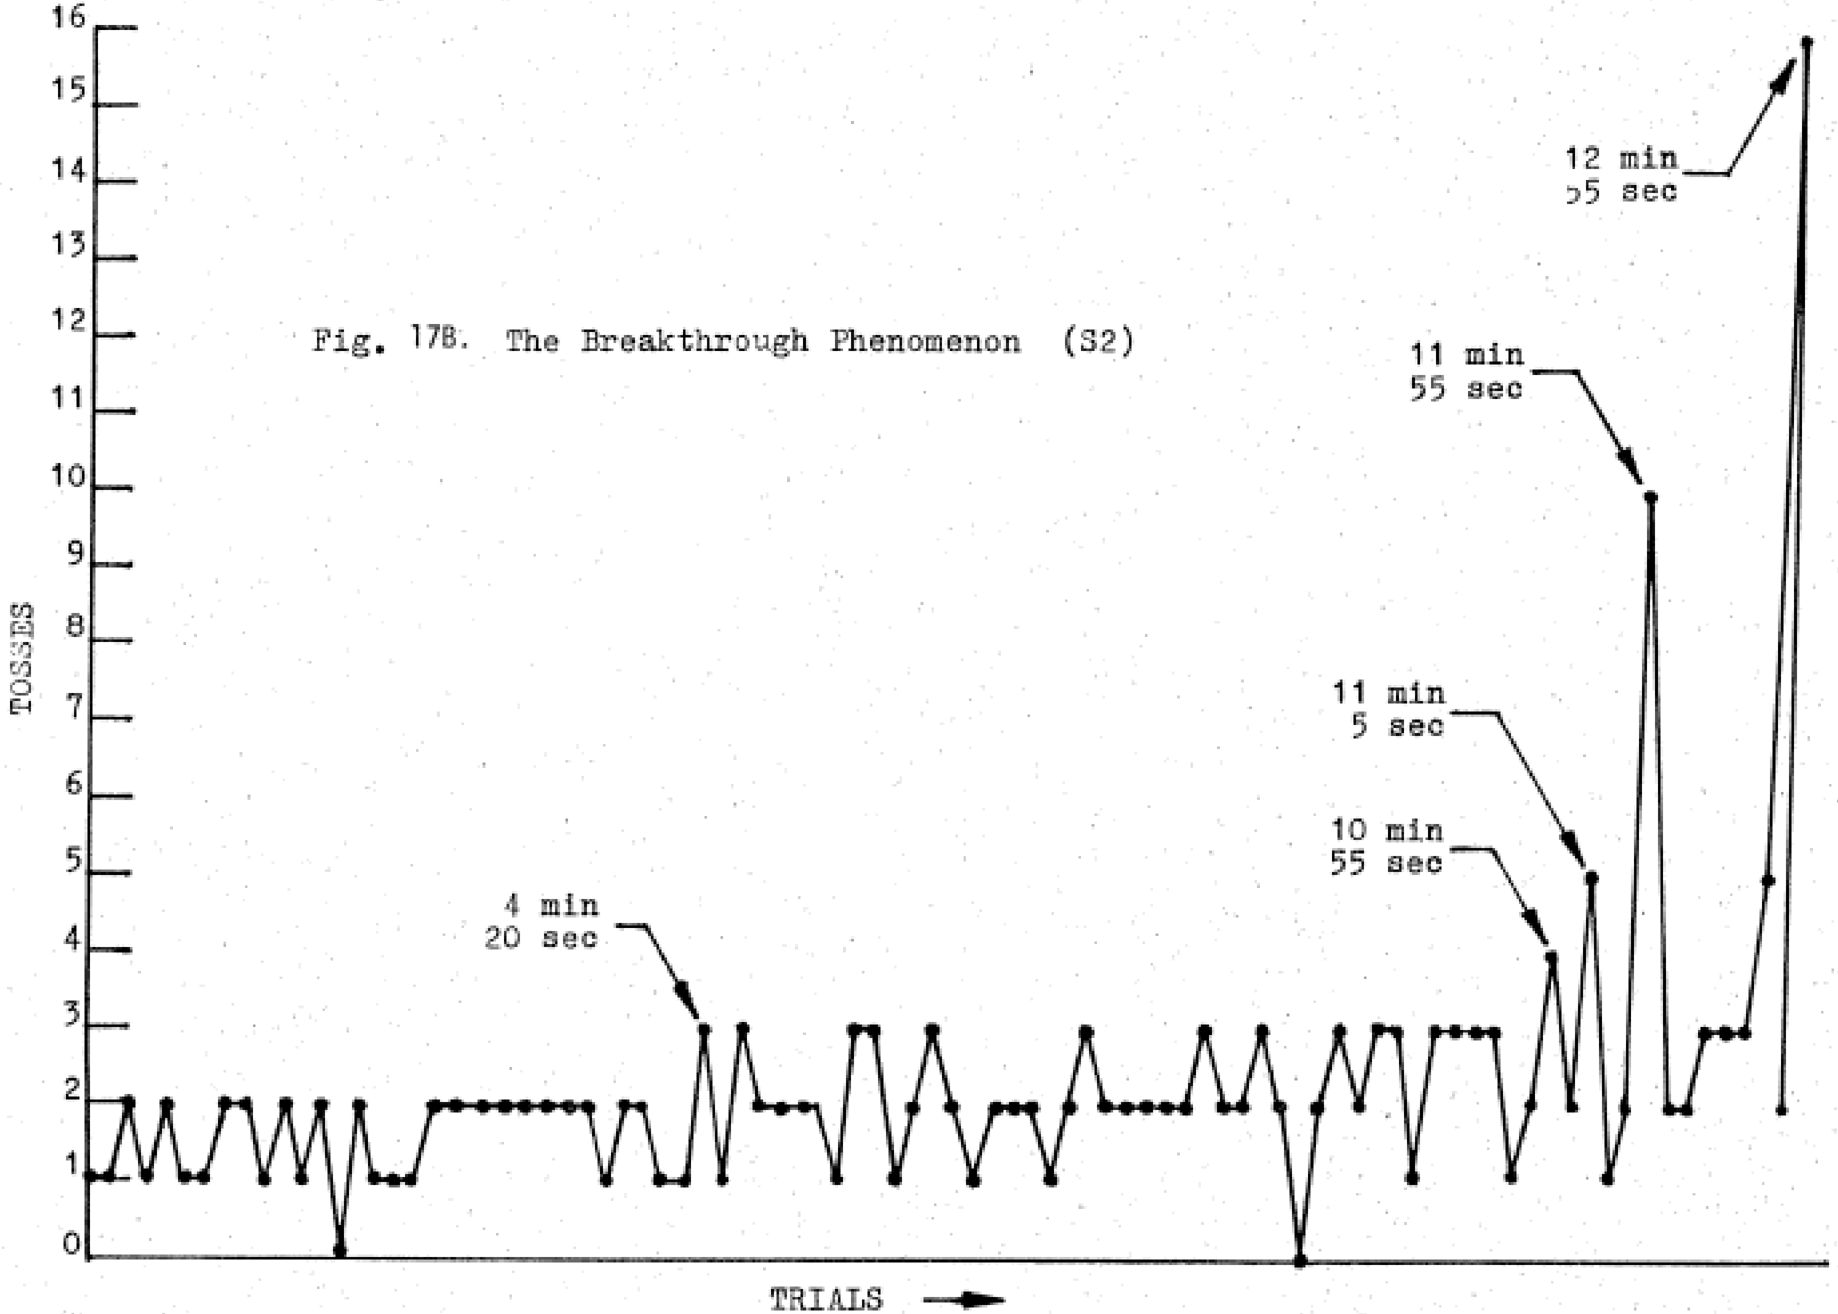

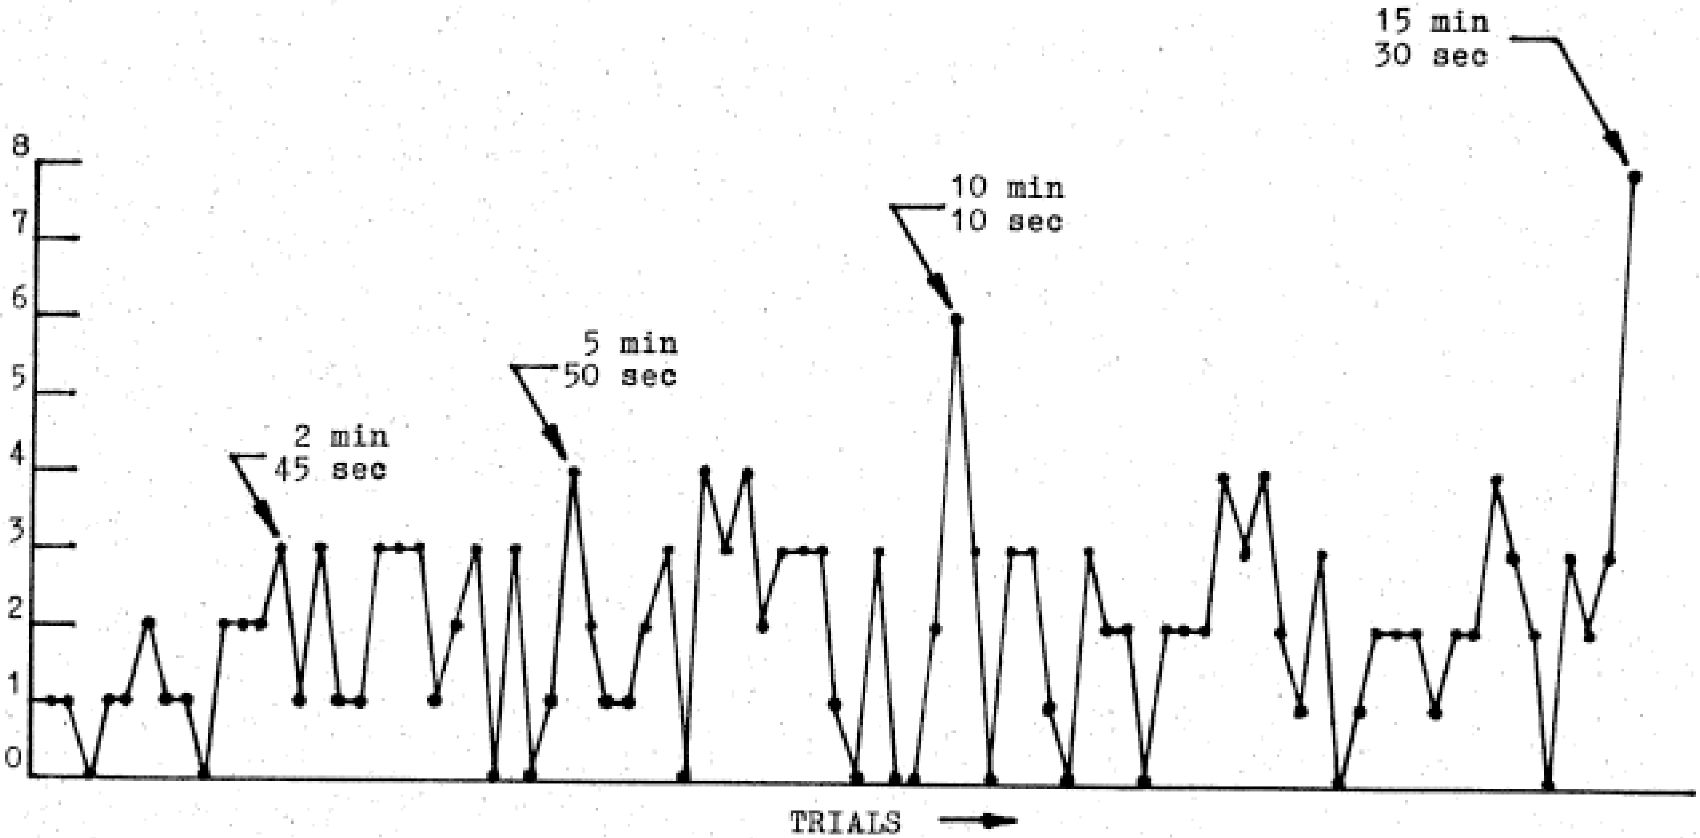

Fig. 17C. The Breakthrough Phenomenon (S5)

It's Hard To Give Yourself Advice

Here's an observation which is true only sometimes and is intriguing both when it does and when it doesn't apply. The question is why is it hard to give yourself advice when learning new physical skills. The most extreme example of this situation is the disastrous "blind trials" approach wherein the subject feverishly tries the task over and over without even taking the trouble to describe his mistakes, let alone advise himself on how to fix them.

> In many cases though, thoughtful analysis has produced a reasonably accurate description of the mistake but it's still difficult to translate that information into corrective action. For example, most people can sucessfully tell themselves to "toss higher" (to gain more time) but apparently it's quite difficult for many people to tell themselves "toss later" or "toss right before _________."

It's also interesting to note that frequently the best advice you can give yourself appears on the surface to have nothing to do with the task at hand but in reality corrects exactly that portion of the movement which was previously erroneously specified. For example in tennis the advice "Serve as though you were trying to throw your racquet away but hold on" frequently helps the novice master the service movement more quickly. In juggling the advice "Practice while facing a wall" is frequently useful for correcting the FORWARD and other bugs.

The theoretical implications of this observation have to do with the mechanisms by which existing movement specifications can be brought to bear on the task of learning new movements. The fact that existing knowledge is relevant to similar task should be obvious and has been demonstrated experimentally by tasks like learning to write with the non-preferred hand. The apparent ability to advise yourself on certain tasks but not on others raises questions about the mechanism by which advice is accomplished and the degree to which it is a symbolic process. Thus the overall theory of physical movement must explain why the symbolic advice "throw this way" (in the serve example) is useful and can be acted upon, but the equally symbolic advice "toss earlier" (in juggling) is only marginally useful.

The Construction of Subroutines

A constant theme which runs throughout the experimental data has to do with the question of how "chunks" of movements are grouped into larger chunks. Computationally this is the issue of sub-procedurization or subroutine construction. It appears to be perhaps the most central issue in designing a theory of physical movement as well as one of the most difficult.

Part of the difficulty lies in ascertaining the real primitives of movement specification and the degree to which the process is symbolic. At least some of the problems posed by this gap in our knowledge can be avoided if we assume that for advanced skills (like juggling) the relevant aspects of the component movements are predominantly symbolic and hence can be treated like ordinary subroutines. Of course this assumption, which may be entirely unwarranted, merely postpones the issue of what happens inside a given subroutine. However it allows us to begin the task of describing the overall learning process from a symbolic point of view.

For example,in juggling,the basic learning process apparently centers around combining the first part of an existing TOSS subroutine

```
TOSS
  DOWNSWEEP
  UPSWEEP
  FLIP
  FOLLOWTHROUGH
```

with the body of an existing CATCH subroutine

```
CATCH
  TRACK
  MOVE-UNDER
  FLEE
  GRASP
```

and inserting MOVE's as appropriate to produce the EXCHANGE procedure

```
EXCHANGE
  TOSS*
  MOVE
  CATCH
  MOVE
```

where TOSS* is TOSS minus the FOLLOW-THROUGH line

Notice that the final version EXCHANGE cannot be a simple linear combination of TOSS, MOVE and CATCH. As previously discussed, the latency times and the real world constraints rule this out. It is certainly true however that, as in Sussman's thesis (see reference 7) the linear approximation subroutine is a good initial model from which to start, i.e. there is a good chance that the debugging process will converge to the desired final structure. The linear model idea also brings to mind the discussion of cascade juggling models in section 4. Sussman and Goldstein (8) also provide clues as to how that convergence process might go in different problem domains. The degree to which those methods apply to the domain of movement depends a great deal on the extent to which movement primitives are symbolic. Issues such as annotation, invocation and patching however, appear to be good starting points for the discussion of whether or not the Goldstein/Sussman theories can be applied to physical procedures. In any event it appears to be true at least at the more advanced levels, that physical movements can successfully be treated as subroutines from a computational point of view.

Summary and Conclusion

In summary, we have seen in considerable detail the extent to which the computational paradigm can be applied to the skill of juggling, a skill which is hopefully representative of most forms of higher-level skill. The discussion initially focused on the external or performance aspects of juggling acquisition. A descriptive language for juggling movements, formulated in bug/fix terminology was developed to the point of exhibiting an extensive taxonomy of cascade juggling bugs.

The latter part of the paper was concerned with establishing the gross outlines of a theory of the internal mechanisms by which movement commands are interpreted and executed. Empirical evidence which was found useful in this regard included an observation about the importance of having a model of the physical task at hand, a study of human ballistic constraints, the unfortunate fact that bugs come in bunches, the very surprising break-through phenomena and the intriguingly difficult task of advising yourself during physical skill acquisition.

This evidence shows that while physical skills have an obvious symbolic component, the precise extent to which the overall theory should be symbolically oriented is not clear. Evidence which supports the symbolic point of view includes the data on models and the successful instances of advice. Evidence which argues against a completely symbolic theory includes the unsuccessful advice cases as well as the "trial-and-error-like" parameter search which is assumed to underlie the breakthrough phenomenon.

The issue of control structure is somewhat more clear-cut. Human and physical constraints make it obvious that a theory of movement must include a control regime which is capable of supporting real-time, parallel events of an anticipatory nature. There must also be a mechanism for describing and modifying event times at the millisecond level. The present experiment yields little information about feedback mechanisms.

The theoretical outline concludes with the observation that the central issue appears to revolve around the problem of how existing movement patterns (subroutines) are combined to form new movements. This statement is confirmed by a subroutine analysis of the juggling learning task. The procedural modification theories of Sussman and Goldstein appear to be partially relevant to physical skill acquisition from that point of view. Hence the issues of invocation, annotation, and self-debugging appear to be good topics for further research. The non-symbolic nature of movement primitives appears to be a central issue in such endeavors. Despite the existence of these problems however, it is felt that the primary conclusions as to the similarities of physical and mental activity and the relevance of the computational paradigm are thoroughly justified.

APPENDIX

Protocol for S1*

1) E-Show me your model of juggling. Make the bags move as though you were juggling.

2) S1-Well let's see. You throw this one (i.e. the left one) in the air. When it reaches its high point you then begin (the transfer operation?)

3) E-Just make them move through the motions.

4) S1-This one goes up (i.e. the left one) and this one goes over (i.e. the right one).
Tosses L→R, passes R→L.

5) E-You have the **model** of juggling that most people have, namely this (demonstrates "shower" juggling). You <u>can</u> juggle that way but it's much harder. I'm going to teach you a <u>new</u> kind of juggling. Let's get some terms. This is a toss (demonstrates a toss). This is a pass (demonstrates a pass). There's no passing in the juggling I'm about to do. It's all tosses (demonstrates cascade juggling).

6) S1-Un huh.

7) E-I'd like for you to do this. One in each hand TOSS R L, TOSS L R. Now notice I don't want <u>simultaneous</u> tosses (demonstrates) and I don't want <u>almost simultaneous</u> (demonstrates). I want a distinct cadence, toss, toss.

8) unintelligible

9) S1-L→R,
R→L, pass, dropped.

10) E-You passed

11) S1-Yeah.

12) L→R "handslip"
R→L wide, forward.

13) L→R fingernail
R→L wide, forward.

14) E-Toss the second one more like that (demonstration)

S1-I'm used to tossing with this hand (left) but not this one (i.e. right).

E-Toss it in a nice parabolic arc.

*In this protocol S1 is the subject, E is the experimenter, L→R is a toss from the left hand to the right hand, ✓ means a satisfactory toss and terms like "handslip", "wide," "forward" the names of specific bugs from the taxonomy, pgs.17-22.

15) R → L ✓ (✓ means satisfactory TOSS).

16) R → L ✓

17) E-Right, now do the exchange.

18) R → L slightly forward, handslip
L → R back, "pinch catch" (i.e. thumb and first two fingers).

19) L → R ✓
R → L slightly wide, slightly forward.

20) E-O.K., quit. That took exactly 2 1/2 minutes. Now do the same thing again but hold the third beanbag in your hand. Just hold it there, don't toss it (demonstration).

21) L → R ✓
R → L low.

22) L → R ✓
R → L ✓

E-O.K..

23) L R back, dropped.
R L very forward.

E-O.K.

24) L → R ✓
R → L forward

S1-I have trouble tossing with this (right) hand.

25) L → R ✓
R → L ✓

E-Please talk out loud.

26) L → R ✓
R → L fingernail

S1-I can really toss well with my left hand but my right hand (isn't too coordinated?).

27) E-Looks good to me. Now what I'd like you to do is when you're at the

stage when the second one's coming down, toss the third one (demonstrates) It goes like this (slow demonstration).

28) L→R slightly forward
R→L very forward, dropped
L→R wide and forward, caught by stepping forward.

E-Almost.

29) L→R ✓
R→L very forward, very wide, stepped forward, dropped
L→R caught directly over shoulder due to previous step.

E-Comment to cameraman.

30) L→R ✓
R→L vertical
L→R very forward, lungecatch

E-Your second toss went straight up and straight down.

S1-Yeah

31) L→R ✓
R→L slightly back, dropped
L→R slightly forward

E-Ah you had it but you dropped the second one.

S1-It's hard for me (to concentrate on catching?).

32) L→R ✓
R→L slightly back, dropped
L→R very forward, lungecatch

E-You're getting a very nice cadence there. Remember that cadence.

33) L→R ✓
R→L slightly back, chestcatch
L→R fingernail

E-O.K., quit. That took you a total of 3 minutes, 35 seconds.

34) Normally we have people stop here because there's about 3 important points in juggling; 3, 10 and 100. But you got there so quickly we'd like to have you go on. Let's see how quickly you get to 10.

E-Explains how to keep going, the apex rule for timing the next toss, and about stopping.

S1-Also it helps to throw it really high so you have time to think about it.

E-Right, keep talking out loud.

35) L→L high, vertical, caught with left hand (unintelligible)

36) L→R very high, back (unintelligible)

37) L→R high back
R→L forward, fingernail

S1-It's like throwing tennis serves.

38) L→R high, slightly back, dropped
R→L

39) L→R ✓
R→L late, caught with R hand.

40) L→ R vertical, chestcatch.

41) E-Notice that it comes and goes. All learning goes like that. Everything goes wrong then it comes back for a couple of times, then . . .

S1-I know how that is.

42) L→R fingernailed
R→L late
L→R back wide, dropped
R→L very late, just rolled out of his hand while catching the third toss (slapdrop?).

E-O.K., you almost got rid of the fourth toss.

43) L→R back, handslip
R→L late, very forward

44) 1) L→R ✓
2) R→L very forward, lungecatch
3) L→R late
4) R→L late, back, leanback catch
5) L→R very forward, lungecatch
6) R→L very forward, lungecatch

7) L→R (not visible)
8) R→L (partly visible) slightly back
9) L→R (not visible)
10) R→L (partly visible) lungecatch bounced loose, chestcatch
11) L→R very back, stepback-catch.

45) E-That took a total of 6 minutes 15 seconds. O.K., let's see what fine timing looks like. Comments on juggling as an unstable (in the mathematical sense) process (advice about "laying" a toss up beside the incoming one and a demonstration).

46) L→R ✓
R→L late, vertical, caught with R
(unintelligible)

47) L→R ✓
R→L late, vertical, caught with R
A little more flip, you don't want to wait so long.

48) L→R ✓
R→L forward, lungecatch

E-You have too much followthrough in your left hand there.

49) L→R slightly back, fingernail
R→L forward, lungecatch

E-Right, much more like that.

50) L→R ✓
R→L forward, lungecatch

51) L→R ✓
R→L late, forward, lungecatch.

52) E-Advice ref. "popping"theory, demonstration.

53) S1-Practices popping (six times)

E-O.K. Hit it. You've already done 11.

54) L→R ✓
R→L forward, lungecatch
(unintelligible)

55) L→R ✓
R→L slightly wide.

56) L→R ✓
R→L wide.

57) L→R ✓
R→L slightly forward, slightly wide
(unintelligible).

58) L→R ✓
R→L✓
(unintelligible)

59) L→R ✓
R→L ✓

E-O.K.. Go ahead. You've done 12 already.

60) L→R vertical, caught with L.

61) L→R vertical caught with L

S1-O.K..

62) L→R low, handslip
R→L later vertical caught with R.

S1-unintelligible.

63) L→R late
R→L late
L→R slapthrow, both dropped

E- Don't slap it. Toss it higher.

64) L→R✓
R→L back, handslip
L→R vertical, caught with L

S1-(Getting in trouble because I'm not getting them out far enough??)

65) L→R fingernail
R→L ✓
L→R very low, very forward, lungecatch

E- You keep getting lower.

66) 1) L→R ✓
2) R→L late, forward
3) L→R forward, lungecatch
4) R→L back
5) L→R almost vertical, slightly low
6) R→L late
7) L→R
8) R→L late, chestcatch, dropped
9) L→R two handed catch

S1-O.K.. I'm not getting the tosses (unintelligible).

E-That's O.K., you're getting good. Keep doing that and you'll clean it up yourself.

67) 1) L→R ✓
2) R→L forward, lungecatch
3) L→R slightly vertical, handslip
4) R→L late, two handed catch

E-Most people don't use the height enough to cheat and get more time up there.

68) 1) L→R fingernail
2) R→L forward, lungecatch
3) L→R forward, lungecatch
4) R→L very wide, leancatch
5) L→R not visible
6) R→L not visible
7) L→R not visible
(unintelligible).

69) L R vertical.

70) 1) L→R ✓
2) R→L slightly forward and left
3) L→R ✓
4) R→L wide, rotates left to catch it
5) L→R vertical, rotates even further left
6) R→L not visible
7) L→R not visible, apparently slightly vertical
8) R→L forward
9) L→R vertical, forward, lungecatch, dropped
10) R→L forward, lungecatch

E-Looks good to me.

S1-Still moving around too much.

71) 1) L→R ✓
2) R→L forward
3) L→R forward, leancatch
4) R→L forward, lungecatch
5) L→R forward, leancatch
6) R→L forward, handslip
7) L→R forward, lungecatch
8) R→L very back, dropped.

72) E-O.K., you do in fact have some indication of a migration bug. The way you fix a migration bug is face a wall (demonstration, more advice ref. migration bug)

73) (now against wall)
L→R wide, back

74) L→R ✓
R→L passes!

75) L→R wide, fingernail
R→L back, handslip
L→R late, low, hits the wall
(unintelligible)

76) 1) L→R fingernail
2) R→L ✓
3) L→R fingernail
4) R→L wide, leancatch, fingernail
5) L→R forward
6) R→L forward, leancatch
7) L→R back, chestcatch, dropped
8) R→L wide, fingernail
9) L→R slightly vertical

E-The one that got you was the one that came back.

S1-Yep.

77) L→R forward
R→L wide

78) L→R vertical, caught with L

E-They tend to go straight up and down.

79) L→R handslip
R→L vertical, caught with R.

80) L→R ✓
R→L vertical, fingernail, chestcatch, leanback
L→R forward, handslip
R→L leancatch
(unintelligible)

81) L→R fingernail
R→L ✓
L→R forward
R→L handslip
L→R vertical, forward, almost hit the wall
R→L wide, forward, leancatch.

82) L→R ✓
R→L ✓
L→R forward, vertical
R→L slightly wide, handslip
L→R early forward, caught against wall
R→L handslip.

83) E-O.K., quit wasting film on him. That took a total of 11 minutes

S1-I was just getting the hang of it.

References

1) Winograd, T., Understanding Natural Language, Academic Press (1972).

2) Minsky, M., Ed., Semantic Information Processing, M.I.T. Press (1968).

3) Minsky & Papert, Artificial Intelligence, Condon Lectures, Oregon State System of Higher Education, Eugene, Oregon (1974).

4) Carlo, The Juggling Book, Random House (1974).

5) Birdwhistle, Kinesics and Context, U. Penn. Press.

6) Lacy, R., "A Model of Juggling" (unpublished paper), M.I.T. A.I. Laboratory (Oct. 1971).

7) Sussman, G.J., A Computational Model of Skill Acquisition, AI-TR-297, M.I.T. A.I. Laboratory (Sept. 1973).

8) Goldstein, I.P., Understanding Simple Picture Programs, AI-TR-294, M.I.T. A.I. Laboratory (Sept. 1974).