

Copyright © 1977, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

AND NOW FOR SOMETHING COMPLETELY DIFFERENT:

The Texas Instruments SR-52

by

W. Kahan

Memorandum No. UCB/ERL M77/23

6 April 1977

INTERLIBRARY LOAN DEPARTMENT
(PHOTODUPLICATION SECTION)
THE GENERAL LIBRARY
UNIVERSITY OF CALIFORNIA
BERKELEY, CALIFORNIA 94720

AND NOW FOR SOMETHING COMPLETELY DIFFERENT:

The Texas Instruments SR-52

by

W. Kahan

Memorandum No. UCB/ERL M77/23

6 April 1977

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

AND NOW FOR SOMETHING COMPLETELY DIFFERENT:

The Texas Instruments SR-52

by W. Kahan

March 1977

This report responds to repeated requests for explanations of the arithmetic paradoxes perpetrated by the T.I. SR-52. The writer has not been paid for this report, nor aided nor abetted in it by Texas Instruments nor by Hewlett-Packard.

Contents

- §0 So What?
- §1 Games and Puzzles
- §2 Where do the digits go?
- §3 Explanations and excuses
- §4 Explanations do not excuse
- §5 So What?
- §6 Conclusion
- §7 FLASH! The New T.I. SR-52 is Different Again

§0 So What?

Of course, all the anomalies to be described in this report are tiny, even negligible from some points of view. Engineers and scientists accustomed to slide rules cannot get upset over errors in the 7th sig. dec. They might get upset if I showed them grossly and unexpectedly wrong results in their own calculations, and often I alienate them in just that way. But alienation does not enlighten, and moreover such picador tactics are rendered futile by the following theorem:

Anything that can be calculated with a sanitary calculator like the hypothetical SR-52-III to be described in §4 can also be calculated with a grubbier calculator like the actual SR-52.

Alas, this theorem does not say how much more the calculation must cost on the SR-52 than on the SR-52-III; in fact the ratio of costs can be made impressively large only by the choice of suitably complicated examples. Consequently no argument about the relative merits and costs of the SR-52 and SR-52-III can be conclusive unless it includes masses of detail more likely to distract than persuade. Rather than drown the reader in such detail, I shall attempt to convey by example just this message:

One man's Negligible is another man's All.

Our example comes from an advertisement for Texas Instruments' SR-60 on p. 23 of the Eastern edition of the Wall Street Journal, Jan. 18, 1977. The SR-60 is a programmable desk-top calculator with a printer, a large display, 95 buttons on its keyboard, 40 to 100 storage cells, 480 to 1920 program steps, and a 12 sig. dec. arithmetic engine similar to that in the SR-52 or SR-51-II. The price, depending upon the amount of memory purchased, ranges from \$1695 to less than \$5000 at this writing. The SR-60 reads

library programs from magnetic cards, and one of these programs was the subject of the advertisement.

The program calculates price and yield for bonds. The advertisement's example concerned a bond with a face value of \$1000, market price of \$1012.15, semi-annual coupon rate of 5.4% per annum, maturity value (call price) of \$1150 if redeemed (called) on June 1, 1979 (the maturity or call date), to be purchased on March 20, 1977. The calculation, to be based upon a fictitious but conventional 360 day/year calendar, must first determine the purchase price taking into account the accrued interest on the next coupon (June 1, 1977) owed to the seller, and then the yield-to-maturity on the buyer's investment stated as a percentage per annum. Here are the results of the calculations:

Calculator	Purchase Price	Yield-to-Maturity
SR-60 as advertised by T.I.	\$1028.33	10.9837 % p.a.
Me, using various machines	\$1028.50	10.9751 % p.a.
Discrepancy	17¢	.0086 % p.a.

The discrepancy in yield exceeds that allowed by the Securities Industry Association manual "Standard Securities Calculation Methods" by Spence, Graudenz and Lynch, 1973, p. 13:

"Calculations for yield should be accurate to four places after the decimal point, rounding to three...just prior to display of yield."

That accuracy is demanded here to ensure that bond dealers will not mislead their customers when they quote assured yields.

The discrepancy in price, 17 cents in \$1028, may seem negligible to an engineer; but it is comparable with the typical profit of 1/64 of 1%,

i.e. 16 cents, made by the bond dealer when he serves only as a broker, except that a transaction will typically involve not \$1028 but more likely \$10,280,000 worth of bonds and the profit will be about \$1600. (For a description of bond traders' activities see Paul Blustein's article in FORBES, Feb. 1, 1977, pp. 37-42.)

The discrepancy in the bond's purchase price could obliterate the profit, or if it went the other way it would pay for the calculator. Ironically, an experienced bond trader can calculate the correct price for this example by hand on the back of an envelope, so he can see from the advertisement that the SR-60 is not for him. Who else can afford to pay its price? If my calculations are correct, the SR-60 and its bond price/yield program are not yet accurate enough for their intended market, though probably accurate enough for engineers who invest some of their savings in bonds. However, engineers can do the necessary calculations on much smaller and cheaper calculators, as I did.

§1 Games and Puzzles

Here is a keystroke sequence to conjure with on the T.I. SR-50, 50A, 51, 51A, 52 or 56 (but not on the 51-II):

Keystrokes	Responses:	Expected	Displayed	Actual
Switch to D		--	--	Degree Mode
Clear All		0	0	0
INV COS		90	90.	89.999 999 99 ₉₈₇
SIN		1	1.	1.000 000 000 ₀₀₁
INV SIN		90	1.Blinking	Blink previous value
- 1 =		?	Varies	See Text

The difference between the last two columns comes about because the calculators display at most 10 leading digits of their results though they carry 13 significant digits in almost all calculations. The difference between what you see and what you get can sometimes be displayed by subtracting what you see from what you got. (The words "almost" and "sometimes" above will be explained later.)

The error in $\text{INV COS}(0)$ is gratuitous but rarely bothersome. The error in $\text{SIN}(\text{INV COS}(0))$ is a nuisance, especially in a programmable calculator, and doubly so when it is hidden from normal view. The angles x at which $\text{SIN}(x) > 1$ are not all known to me but they include instances like

$$\text{SIN}(89.99995) = 1.000\ 000\ 000_{004}$$

for which I thank Hugh Ross of Poughkeepsie, N.Y. His discovery was made on the SR-51; it could not so easily have been accomplished on the SR-52 or 56 because

$$\text{SIN}(89.99995) - 1 =$$

yields 0. on the SR-52 and 56 but 4×10^{-12} on the SR-51. Perseverance pays; $1 - \text{SIN}(89.99995) = -4 \times 10^{-12}$ on all three machines. (The T.I. SR-51-II lacks the entertaining "features" described above.)

The SR-51 and 52 differ from other calculators in another way; $(\text{SIN}(89.99995))! = 1.$ exactly on the SR-51 and 52 but blinks on the SR-50, $(\text{COS}(5.814886783 \times 10^{-4}))! = 1.$ exactly on the SR-51 and 52 but blinks on the SR-50, but $(\text{COS}(5.814886784 \times 10^{-4}))!$ blinks on all of them. The reason is perfectly clear:

$$\text{SIN}(89.99995) = 1.000000000_{0,04} \text{ displays as } 1.$$

$$\text{COS}(5.814886783 \times 10^{-4}) = .999999999_{5,00} \text{ displays as } 1.$$

$$\text{COS}(5.814886784 \times 10^{-4}) = .999999999_{4,90} \text{ displays as } .999999999$$

The last displayed value is the only one not an integer, and only non-negative integers will be accepted by the T.I. calculators' factorial function $x!$.

The later machines' factorial $x!$ differs from the earlier SR-50's by acting not on x but on the 10 sig. dec. rounded value of x in the display. In this respect $x!$ differs too from all the other functions on the SR-51 and 52. I do not know why a similar fix-up was not applied to the inverse trigonometric functions. Neither do I know why a different fix-up for $x!$, which could have been extended smoothly to provide the gamma function $\Gamma(1+x)$ at non-integer values x , was rejected. Some other calculators (e.g. Commodore SR5190R) do provide the gamma function; it is useful to statisticians among others.

It is not easy to draw the right conclusions from tests performed upon calculators. The next few examples were suggested by Fred Powell of Niagara Falls, N.Y., who obtained the values marked (*) below on his SR-56. He enjoyed "wry amusement" at finding 2^3 "exactly equal to 8", when using the y^x key, despite my contrary allegations; and he was not dissatisfied when 55^4 fell only 0.0001 short of its true 7-digit integer value 9 150 625. I enjoyed mixed feelings when I obtained the following results on the SR-52 and 56, and I was not satisfied when calculations with small integers which should have delivered small integers (small compared with the machine's capacity) delivered something else; and I regret that Mr. Powell's calculator did not answer truly when he asked it whether 2^3 exactly equalled 8.

$2^3 - 8 = 0. \quad (*)$	$8 - 2^3 = -1 \times 10^{-12}$
$(55^2)^2 - 55^4 = 0.000092$	$55^4 - (55^2)^2 = -0.0001 \quad (*)$
$(2^4)^8 - (2^8)^4 = -.215$	$(2^8)^4 - (2^4)^8 = .21 \quad (**)$

Perhaps he will be pleased to learn that 55^4 is slightly better than he thought. Others, who purchased what they thought was a 12 or 13 digit T.I. calculator instead of a competitor's 10 digit calculator so that they could apply theorems like Fermat's to a wider range of integers, will entertain second thoughts. (Fermat's theorem says of integers x and p that $(x^{p-1}-1)/p$ is an integer whenever x/p is not an integer and p is prime, and is useful for testing whether a large integer p is prime.) They may have believed that no calculator whose error is smaller than 1 unit in its last digit displayed could miscalculate integers unless those integers exceeded the capacity of the calculator's registers, but this truism is true only if the last digit displayed really is the last digit.

Another cherished belief is that any function $f(x)$ should take the same value on different occasions if calculated for the same value of the argument x ; in other words, $f(x) - f(x) = 0$ even if the calculated value of $f(x)$ on both occasions is wrong. Mr. Arthur Schlang of Melville, N.Y., discovered that this equation is hardly ever satisfied by non-trivial functions f on the SR-52; for instance he found

$$\sqrt{125} - \sqrt{125} = -9 \times 10^{-11} .$$

He also found a cure; "all difference problems disappear on the SR-52 if the subtrahend is post multiplied by unity thus:"

* and ** The T.I. SR-51-II produces these marked values but not the others.

$$f(x) - f(x) \times 1 = 0 \quad (\text{but not necessarily } f(x) - 1 \times f(x)) .$$

Is this cure preferable to the disease? Another cure is the following:

Calculate $f(x)$, without reference to cell #19, say.

ST θ 19 to save it in cell #19.

Re-Calculate $f(x)$ without reference to cell #19.

INV SUM 19 (subtract $f(x)$ from cell #19).

RCL 19 to display $f(x) - f(x) = 0$ every time.

This cure works only on the SR-52 and 56. The SR-51 has a similar disease which is not cured but contracted via operations like those above; if $f(x)$ is calculated and stored on the SR-51, then recalculated and subtracted from its stored value, a non-zero difference will often be displayed. The SR-51-II appears not to have either disease.

§2 Where do the digits go?

By now the reader may have figured out what is going on. The SR-50, 51, 52 and 56 return 13 significant digits for every arithmetic operation and elementary function, though the 13 figures might not all be correct. The SR-51 keeps only 12 of those 13 figures when it stores into one of its 3 memory cells; the last digit is chopped off. The SR-52 and 56 keep all 13 digits when they store or do arithmetic into storage; but they keep only 12 in their "hierarchy" stacks into which numbers are pushed after any one of the operator keys $+$, $-$, \times , \div , y^x , $y^{1/x}$ or $($ is depressed. The 13th digit appears to be struck off to make a place where the operator-key may be remembered, but I am not sure of this. Consequently $f(x) - f(x)$ becomes $f(x)_{12} - f(x)_{13}$ where the subscripts denote the number of significant decimals retained. Similarly $A \times B$ becomes $A_{12} \times B_{13}$, and may be different from

$B_{12} \times A_{13}$. Now we can understand Mr. Schlang's cure; in replacing

$$f(x) - f(x) \neq 0 \text{ by } f(x) - f(x) \times 1 = 0$$

he has really replaced

$$f(x)_{12} - f(x)_{13} \neq 0 \text{ by } f(x)_{12} - f(x)_{12} \times 1_{13} = 0 .$$

The second cure, using 13 digit storage cells, yields

$$f(x)_{13} - f(x)_{13} = 0 .$$

But something else can still go wrong.

When we use 13 digit arithmetic into storage on the SR-52 or 56 we still find that $A \times B - B \times A \neq 0$ for many choices of A and B despite that we expect to calculate

$$(A_{13} \times B_{13})_{13} - (B_{13} \times A_{13})_{13} = 0 .$$

For instance, these calculators and the SR-51 say

$$3.16227\ 7660_{169} \times 3.16227\ 7660_{170} = 10.0000\ 0000_{000} \text{ but}$$

$$3.16227\ 7660_{170} \times 3.16227\ 7660_{169} = 9.9999\ 9999_{978} ;$$

the last result seems inexplicable considering that both factors exceed $\sqrt{10} = 3.16227\ 7660_{168379} \dots$ and their true product is $10.0000\ 0000_{000708} \dots$.

Non-commutative multiplication on these machines appears to be caused by a lack of guard digits and consequent premature loss of information; for instance, the last (13th) digit of the second factor sometimes fails to contribute to the product. Since the discrepancies affect only the 12th and 13th significant decimals of products, they are "out of sight, out of mind" on a calculator that displays only the first ten.

Addition and subtraction suffer more seriously from lack of a guard digit because it prevents the difference between nearby numbers from being determined correctly unless they share the same exponent of 10 in scientific notation. For instance, the SR-51, 52 and 56 calculate $3 \times \frac{1}{3} = 9.9999\ 99999_{999} \times 10^{-1}$ correctly to 13 sig. dec., though they display its value rounded to 10 sig. dec. as 1×10^0 ; but they calculate

$$1 - 3 \times \frac{1}{3} = 10^{-12} \quad \text{though} \quad .5 + (.5 - 3 \times \frac{1}{3}) = 10^{-13} .$$

The two calculations differ only in their vulnerability to a missing guard digit:

1.0000 0000 ₀₀₀		.50000 00000 ₀₀₀
- 0.9999 99999 ₉₉₉	← This 9 drops off for lack of a guard digit.	- .99999 99999 ₉₉₉
0.0000 00000 ₀₀₁		- .49999 99999 ₉₉₉
= 10^{-12}		+ .50000 00000 ₀₀₀
	$10^{-13} =$.00000 00000 ₀₀₁

Texas Instruments no longer advertises the SR-51, 52 or 56 as 13 digit calculators; like the SR-51-II they are advertised as 12-figure calculators though the earlier calculators are described as "carrying" 13 figures. The newer SR-51-II appears to calculate in the same 13-figure way as its predecessors but chops the 13th figure off each result before presenting it. For instance, the SR-51-II delivers $1 - 10^{-12}$ instead of $1 - 10^{-13}$ when it calculates $3 \times \frac{1}{3}$, and says $\text{SIN}(89.99995) = 1$ instead of $1 + 4 \times 10^{-12}$. Consequently it behaves like a 12 figure calculator equipped with a 13th guard digit, and lacks many of the entertaining features of its predecessors; differences are calculated correct to within a unit in their last (12th) significant digit, and multiplication commutes as it should. These calculators

illustrate a principle that many people find paradoxical at first; 12 figures plus a guard digit are preferable to 13 figures without a guard digit. Merely to carry 13 figures is not enough; the way that the 13th is carried is significant too. The SR-51, 52 and 56 carry their 13th figures in ways which, rather than enhance the 12th, cause anomalies that tend to undermine confidence. What other surprises lurk within those calculators?

§3 Explanations and excuses

The French have a word for it: "Tout comprendre est tout pardonner." Knowing about missing guard digits and about the peculiar sub-expression stacks in the SR-52 and 56, we understand how the phenomena described must occur. To know why they occur we must appreciate the design philosophy upon which all the T.I. machines appear to be based. It is the philosophy of the

"Backward Error-Analysis":

Suppose we wish to compute $f(x)$ but are obliged to accept instead an approximation $F(x)$. We might hope that $F(x) = \text{almost } f(x)$ in the sense that $|F(x) - f(x)| \leq \epsilon |f(x)|$ for all pertinent x and some tiny positive ϵ like $\epsilon = 10^{-12}$. Should this hope be confirmed by an error-analysis, we would say that " $F(x)$ submits to a satisfactory Forward Error-Analysis" to distinguish so favourable an approximation from what occurs more often in practice. Most people think in terms of a forward error-analysis because it is easier to interpret and seems more natural, and because they are accustomed to seeing it work for products, quotients, square roots and sums, and sometimes for differences (there is some confusion here). But complicated or lengthy calculations rarely behave so well. Usually, if $F(x)$ is at all an acceptable approximation, the best that can be proved is that $F(x) = \text{almost } f(\text{almost } x)$ in the sense that $|F(x) - f(y)| \leq \epsilon |f(y)|$ for

some unknown $y = y(x)$ which satisfies $|y-x| \leq \epsilon|x|$, and ϵ is some tolerably tiny number. In this case, when $F(x) = \text{almost } f(\text{almost } x)$ we say that "F(x) submits to a satisfactory Backward Error-Analysis" even though, as sometimes happens, $F(x)$ might be nowhere near $f(x)$.

The thought that F might be nowhere near f and still be called a "satisfactory" approximation must jar most sensibilities. Moreover, engineers and scientists must be perplexed when continuous functions f appear to violate a definition of continuity that insists $f(\text{almost } x) = \text{almost } f(x)$. Indeed, engineers who deal with noise in linear systems regard the last equation as a confirmation of their ability to refer noise indiscriminately either to an input or an output. But non-linear systems and many computations do not honour the quantitative version of that equation wherein the word "almost" is defined to mean "within a prescribed tiny relative uncertainty ϵ ". For instance, take $\epsilon = 10^{-12}$ and $f(x) = \tan(x)$ for very large radian $x > 10^{50}$. Changing x in its 13th sig. dec. changes x absolutely by at least 10^{38} radians and would almost surely change $\tan(x)$ utterly; consequently $\tan(\text{almost } x) \neq \text{almost } \tan(x)$ when x is huge. And yet we can hardly expect a calculator to do much better than provide $\text{TAN}(x) = \tan(\text{almost } x)$ because the only practical way known to compute $\tan(x)$ requires first that a suitable integer $n \doteq x/\pi$ be calculated and then $\tan(x) = \tan(x-n\pi)$ will be computed from a moderate size argument $(x-n\pi)$ in the range $-\pi/2 < (x-n\pi) < \pi/2$. However, if only 13 sig. dec. are used to approximate π and to perform arithmetic, the calculated value of $(x-n\pi)$ will be wrong by an amount comparable with $10^{-13}|x|$ whenever $n \geq 1$. Hence $\tan(\text{slightly wrong}(x-n\pi)) = \tan(\text{almost } x) \neq \text{almost } \tan(x)$ whenever $|x|$ is huge. Therefore we should not expect a computer's TAN to be much better for radian arguments x than $\text{TAN}(x) = \tan(\text{almost } x)$, although by

devious means unsuited to discussion here better performance is possible. On the other hand, if the angle x is represented in degrees the foregoing argument is inapplicable because 180, unlike π , is representable precisely in computers; hence we might expect $\text{TAN}(x) = \text{almost tan}(x)$ for all degree arguments x including odd multiples of 90° where $\text{tan} = \pm\infty$.

Let us now compare actuality with the expectations aroused by the foregoing discussion. First we calculate $\text{tan}(x)$ for a large radian argument $x = 52174$:

Calculators	Calculated TAN(52174 radians)
Texas Instruments SR-51,52,56	-181683.7387 ₆₆₉ = tan(52173.99999 _{999 6561...})
Hewlett-Packard HP 67,91,97,27	-181683.5708 = tan(52173.99999 _{999 6566...})
True tan(52174)	-181570.29570...

Note that these calculators err in the 4th significant decimal of $\text{tan}(x)$, but they produce a value $\text{tan}(\text{almost } x)$ where "almost x " differs from x only after its 13th significant decimal. Scant cause for complaint here.

Our next calculation compares $\text{tan}(90+x \text{ degrees})$ with $\text{tan}(90-x \text{ degrees})$ for tiny $x = 10^{-8}$; we should get tangents differing only in sign.

Calculators	TAN(90.0000 0001°) TAN(89.9999 9999°)
T.I. SR-51,52,56	-57515 65971. ₂₆₃ = tan(90.0000 0000 _{996 177...} °) 57178 91696. ₉₀₄ = tan(89.9999 9998 _{997 956...} °)
H-P HP-27,67,91,97	-57295 77951. correct to 10 sig. dec. 57295 77951.
Monroe 326	-57295 77951. ₃₀₆ correct to 12 sig. dec. 57295 77951. ₃₀₆

The failure of $\text{TAN}(90+x^\circ) = -\text{TAN}(90-x^\circ)$ in the T.I. machines' third sig. dec. is disconcerting first because it is unexpected and second because the failure does not occur on those other calculators whose calculated tangents are in error by less than one unit in the last place displayed. The fact that still T.I.'s $\text{TAN}(x^\circ) = \tan(x^\circ$ to within 12 sig. dec.) is less satisfying now than before, especially when we find that the performances tabulated above are typical for trigonometric functions in degrees on the machines cited.

So far, all the T.I. machines' arithmetic anomalies exhibited in this report are consistent with a Backward Error Analysis,

$$\text{calculated } F(x) = \text{almost } f(\text{almost } x) ,$$

with "to 12 sig. dec." as the interpretation of "almost". The prefixed "almost" is not contradicted when $\text{SIN}(89.99995^\circ)$ exceeds 1 by only 0.4 units in the 12th sig. dec. The parenthesized "almost" is compatible with the loss of a 13th digit in the sub-expression ("hierarchy") stack that causes anomalous results to appear for

$$\text{calculated } (x-y) = (\text{almost } x) - (\text{almost } y)$$

$$\neq - \text{calculated } (y-x) = (\text{almost } y) - (\text{almost } x) ;$$

just remember that each of the four "almost"s can differ from the others (i.e. $\text{almost } x \neq \text{almost } x$) even though all "almost"s are almost the same (i.e. they agree with 1 to 12 sig. dec.). Similarly compatible with backward error analysis are the consequences of missing guard digits that allow

$$\text{calculated } (x \cdot y) = (\text{almost } x) \cdot (\text{almost } y)$$

$$\neq \text{calculated } (y \cdot x) = (\text{almost } y) \cdot (\text{almost } x)$$

and

which is the answer correct to 10 sig. dec. calculated by the Monroe 326 and the HP-22, 27, 67, 91 and 97. Then $Y^X = \text{EXP}(X \cdot \text{LN}(Y))$ is affected adversely in consequence. Unfortunately, the anomalous logarithms and exponentials persist in the newer SR-51-II.

§4 Explanations do not excuse

Each anomaly seems less objectionable after it has been explained by a backward error analysis, but therein lies little satisfaction. The French expression cited above,

"Tout comprendre est tout pardonner",

is wrong. The explanations do not excuse the anomalies in their aggregate. We feel like the players in a game, whose rules we have come to accept as inevitable although we do not yet fully understand them, who have just been told that henceforth we must play by numerous new rules which will vary slightly from game to game in ways to be determined by an anonymous umpire who declines to announce them in advance because, he says, the variations are almost inconsequential.

The anomalies, though sanctioned by backward error analysis, are objectionable to the extent that they are unnecessary. Except for the trigonometric functions of large radian arguments, all the functions on today's calculators can be delivered very nearly correctly rounded (to within say .51 units in the last figure carried) at no significant disadvantage in price or performance relative to today's calculators. The advantages of a calculator free from unavoidable anomalies are expressible in three words,

economy of thought,

but they are worth spelling out less cryptically:

Let us examine the advantages over the T.I. SR-52 of a hypothetical SR-52-III differing from its predecessor only in its freedom from the arithmetic anomalies discussed above; we assume that the SR-52-III would deliver correctly rounded 12 sig. dec. values for the algebraic functions (+, -, ×, ÷, $1/x$, \sqrt{x} , x^2) and values for the other elementary functions (logarithmic, exponential, y^x , $y^{1/x}$, and trigonometric) in error by less than .51 units in the 12th sig. dec. (except for an additional error in large radian arguments of trigonometric functions). The first advantage of the SR-52-III is that its machinations are easier to understand because they are free from distracting surprises. Next we find that programs are easier to devise and considerably shorter in many cases, for instance to calculate hyperbolic functions (sinh, cosh, tanh and their inverses) and financial functions like $((1+x)^n - 1)/x$ to at least 10 sig. dec. Also, roundoff noise causes less confusion in programs to solve nonlinear equations, evaluate integrals numerically, solve differential equations and perform statistical analyses of various kinds of data. These programs can be shared and modified with less risk of malfunction. Because programs are shorter, they are more transparent and contain fewer blunders. In short, we get better answers sooner, more often, and at less cost. And when we get wrong answers (nobody is infallible) the causes are easier to find because they lie in what we have done instead of in what the calculator has done to us. Finally, a subtle but invaluable advantage emerges; when we must decide for sure whether our results are accurate enough for their purpose, we can perform the auxiliary calculations that constitute an error analysis with confidence that our vision of the source of uncertainty is both valid and simple enough to support useful calculation. That no-man's-land between the results known to be accurate and those known to be inaccurate is appreciably narrower for the hypothetical SR-52-III than for the actual SR-52.

To people inexperienced at long computations the foregoing claims must seem like an awful lot to base upon so little. Lacking the hypothetical SR-52-III, I cannot easily prove my claims. But I can supply evidence of a negative character; I can demonstrate surprisingly wrong answers, much more wrong than could be explained by backward error analysis, on T.I. calculators. Some of these wrong answers are described elsewhere*, so I shall add just one more here.

The hyperbolic functions (\sinh , \cosh , \tanh and their inverses) have been provided on the T.I. SR-50, 50A, 51, 51A and 51-II calculators directly, and indirectly in software on the SR-52 and 56. The SR-50 and 50A yielded $\sinh(x)$ and $\sinh^{-1}(x)$ with fairly large relative error for tiny arguments x though the absolute error was small, about 10^{-12} . A revised procedure was incorporated into the SR-51, 51A and 51-II to yield full relative accuracy (12 sig. dec.) for all x . But the algorithm used for $\sinh^{-1}(x)$ suffers from roundoff in intermediate results in a way not anticipated by its designer. Consequently arguments x exist at which the relative error in $\sinh^{-1}(x)$ is almost 2×10^{-7} rather than 10^{-11} ; for instance $\sinh^{-1}(1.000000099 \times 10^{-5})$ is calculated as $.9999999 \times 10^{-5}$ instead of $1.000000099 \times 10^{-5}$. No 12 sig. dec. backward error analysis of the function $f(x) = \sinh^{-1}(x)$ can excuse this error. It would not have occurred if the calculator's designer had provided a logarithm function correct to 12 sig. dec. (in the sense of forward error analysis) and had used it to calculate \sinh^{-1} . Instead he was obliged to patch together at least two different algorithms for $\sinh^{-1}(x)$, choosing one or the other in accordance with x 's magnitude; the patches do not completely cover the hole.

*"Can you count on your calculator?" by W. Kahan and B.N. Parlett, March, 1977.

The software provided by T.I. to calculate $\sinh(x)$ and $\sinh^{-1}(x)$ on the SR-52 suffers from absolute errors of the order of 10^{-12} , but these are relatively large relative errors when x is tiny, and adversely affect calculations of $\sinh(x)/x$ and $\sinh^{-1}(x)/x$. For instance some calculated values of $\sinh^{-1}(x)/x$ which should round to 1 turn out to be

$$\begin{aligned} 0 & \text{ for } 0 < x \leq .99999\ 99999 \times 10^{-12} \text{ and} \\ 1.0000002 & \text{ for } x = 5 \times 10^{-7} ; \end{aligned}$$

the last value violates the well-known inequality

$$0 < \sinh^{-1}(x)/x \leq 1 .$$

These errors are unnecessary artifacts generated by the standard formula

$$\sinh^{-1}(x) = \ln(x + \sqrt{1+x^2})$$

partly because of rounding errors committed during the calculation of $(x + \sqrt{1+x^2})$ and partly because T.I.'s logarithm function calculates $\text{LN}(x) = \ln(\text{almost } x)$. In particular $\text{LN}(1 + 5 \times 10^{-7}) = 5.000001 \times 10^{-7}$ explains why the inequality above was violated. The rounding errors could be avoided via the following procedure:

$$\sinh^{-1}(-x) = -\sinh^{-1}(x) \text{ so assume } x \geq 0 .$$

$$v = x + x^2 / (1 + \sqrt{1+x^2}) ; \quad u = 1 + v .$$

$$\text{If } u = 1 \text{ then } \sinh^{-1}(x) = v ; \text{ otherwise}$$

$$\sinh^{-1}(x) = v \ln(u) / (u-1) .$$

This procedure gives satisfactory results (in the absence of over/underflow) on calculators like the Monroe 326 and HP-22, 27, 67, 91 and 97 which have logarithms that yield $\text{LN}(x) = \text{almost } \ln(x)$; but the procedure is unusable

on the T.I. SR-52 and 56 because their logarithms are not so good. Instead, power series or continued fractions must be developed for the T.I. machines at a significantly greater cost in analysis, programming and performance. And this is just what I claimed; economy of thought has been undermined by a very tiny anomaly in T.I.'s logarithm procedure.

The foregoing procedure for $\sinh^{-1}(x)$ also serves to expose and unravel a paradoxical aspect of backward error analysis. Let V denote the calculated value of v and U the calculated value of u . It is not hard to see that despite roundoff both $V = \text{almost } v$ and $U = \text{almost } u$. Then

$$\text{LN}(U) = \text{almost } \ln(\text{almost } U) = \text{almost } \ln(\text{almost almost } u)$$

is calculated on the T.I. machines, and here we find scant reason to object to the extra uncertainty generated by one more "almost" than would have been generated by the other calculators'

$$\text{LN}(U) = \text{almost } \ln(U) = \text{almost } \ln(\text{almost } u) .$$

How much more uncertain than $(\text{almost } u)$ is $(\text{almost almost } u)$? Surely the difference is negligible?

The different uncertainties would indeed differ negligibly in an isolated calculation of $\text{LN}(U)$. But the error in the argument U of $\text{LN}(U)$ is correlated with the error in the expression $U - 1$. On the other calculators we find that, since $1 < U = \text{almost } u$,

$$\text{Computed } \left(\frac{\text{LN}(U)}{U-1} \right) = (\text{almost})^3 \frac{\ln(U)}{U-1} = (\text{almost})^4 \frac{\ln(u)}{u-1}$$

because $\ln(u)/(u-1)$ is a very tame function; but on the T.I. calculators

$$\text{Computed } \left(\frac{\text{LN}(U)}{U-1} \right) = (\text{almost})^3 \frac{\ln(\text{almost almost } u)}{\text{almost } u - 1}$$

can be nowhere near $\ln(u)/(u-1)$ when u is very near 1. Thus does T.I.'s logarithm preclude exploitation of correlated errors despite satisfying a backward error analysis. In many long computations, correlated errors are not exceptional.

§5 So What?

Please re-read §0 before proceeding to §6.

§6 Conclusion

Before I received letters concerning the SR-52 and what I say about it, I used to believe that no mathematical phenomenon could generate acrimony. I believe that still; the phenomena described above are not mathematical. They are psychological, social and economic phenomena.

History's religious wars suggest that men relinquish only reluctantly those rare opportunities to impose their prejudices upon others. Is this why prejudices about what is "negligible" are imposed, sometimes unwittingly, by the designers of calculators upon their customers? If so, I should try to impose my prejudices upon these designers, and then calculators' errors would be better than negligible; they could hardly be diminished. Also negligible, to those who understand the necessary mathematical technology, would be the increased cost of producing such exemplary calculators. Alas, true enlightenment cannot be achieved by mere religious conversion.

The dominant forces in calculator design appear to be market forces. So long as people who purchase calculators tend to ignorance of or indifference to the ways calculators guard or fail to guard their accuracy, manufacturers will sense no incentive to maintain accuracy; instead they will strive for "features" that seem to stimulate sales. Doing so, each manufacturer is

merely acting out his chosen rôle, which is to produce for profit what people do not say they do not want.

The designers and manufacturers of the T.I. SR-50, 51, 52, 56 series do deserve two accolades. First, they are learning; the SR-51-II has appreciably fewer arithmetic anomalies than its predecessors. Second, they deliver astonishing value for money. So low are their calculators' prices that a professional engineer who dissipates more than a morning unravelling an enigma like those explained in this report will have spent time worth more than the price of his calculator.

§7 FLASH! The New T.I. SR-52 is Different Again (April 15, 1977)

The new model differs from the old in only one external respect that I have noticed; its serial number is much bigger (in the millions) and stamped into the back of the black plastic case instead of being merely printed on the label. Internally there are significant improvements. As far as I can tell from a few minutes' inspection, the new SR-52 uses the same arithmetic as the SR-51-II; i.e., each 13 figure result gotten on the old SR-52 has its last digit chopped off before being delivered by the new SR-52. Consequently the new SR-52's arithmetic appears to be 12 sig. dec. chopped arithmetic with a guard digit, which is must less prone to pathology than before. Multiplication and addition appear to be commutative again, and $|\sin(x)| \leq 1$. But logarithms and exponentials and trigonometric functions still suffer from the anomalies bequeathed by indiscriminate backward error analyses (see §3 and §4 above). In fact, some of the anomalies have been worsened to 12 sig. dec. For instance, consider an example once widely cited by T.I. in some of its promotional literature for the SR-50 and 51:

Enter your telephone number (including the area code), a ten digit integer, into the calculator. Now press the LN key, and then EXP (or INV LN), to calculate $\exp(\ln(\text{your telephone number}))$. Your telephone number is once again displayed unaltered. But if you did this on some competitors' calculators, your telephone number might be changed. This shows that the T.I. calculators are more accurate than those competitors'. Or does it?

To be precise, let us observe that we are asked to calculate

$$f(x) = \exp_r(\ln_r(x)) \quad \text{when } x = \text{your telephone number,}$$

where the subscript "r" means "rounded or chopped in whatever way is built into the calculator". On the various calculators that subscript "r" has different meanings:

old T.I. SR-52	... produce 13 sig. dec. (not all correct)
new T.I. SR-52, SR-51-II	... keep only the first 12 of those 13 sig. dec.
Monroe 326	... produce 13 sig. dec. (almost all correct)
Commodore SR-4148	... produce 12 sig. dec. (not all correct)
Older HP calculators	... produce 10 sig. dec. (not all correct)
Newer HP calculators	... produce 10 sig. dec. nearly correctly rounded

Except for the Monroe 326, which can display up to 12 sig. dec., all these calculators display at most 10 sig. dec.

Now let us calculate $f(x)$ on various calculators and for various values of x ; specifically, we start with

$$x_0 = 919\ 555\ 1212 \quad (\text{INFORMATION's telephone no. in Raleigh, N.C.})$$

and then $x_i = f(x_{i-1})$ for $i = 1, 2, 3, \dots$. This peculiar choice of values x is motivated by a theorem which says that when r means "round to n sig. dec." for any fixed n , say $n = 10$ or $n = 12$, then $x_1 = x_2 = x_3 = \dots$

even though possibly $x_1 \neq x_0$. The table below displays only $x_i - 919\,555\,1000$ to save space.

$x_i - 919\,555\,1000$

i	old SR-52	SR-51-II	Monroe 326	Commodore SR-4148	Older HP	Newer HP
0	212.	212.	212.	212.	212.	212.
1	211.968	211.69	211.971	211.13	210.	211.
2	211.876	210.77	211.971	210.21	210.	211.
3	211.784	209.85	211.971	209.28	210.	211.
⋮	⋮	⋮	⋮	⋮	⋮	⋮
6	211.508	207.09	⋮	206.53	⋮	⋮
7	211.416	206.17	⋮	205.61	⋮	⋮
⋮	⋮	⋮	⋮	⋮	⋮	⋮
22	210.129	⋮	⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮	⋮	⋮	⋮
45	⋮	174.91	⋮	170.66	⋮	⋮
⋮	⋮	⋮	⋮	⋮	⋮	⋮
∞	210.129	174.91	211.971	?	210.	211.

The interpretation of this table requires some care. Note that the displayed values are obtained by rounding off the small-type digits; for instance, on the old SR-52 we would see displayed $x_0 = x_1 = \dots = x_5 = x_6 = 919\,555\,1212 \neq x_7 = 919\,555\,1211$, on the SR-51-II we see $x_0 = x_1 \neq x_2 = 919\,555\,1211$, and so on. Obviously the 12 and 13 figure calculators are more accurate initially than the 10 figure calculators. But the 10 figure calculators and the Monroe 326 are the only ones which do not violate the theorem for this value x_0 . For other values of x_0 only the new HP calculators do not violate the theorem.

Which is more important; the "accuracy" of a result measured in correct leading sig. dec., or the satisfaction of various theorems about the relations satisfied by correctly rounded calculations? Don't answer too soon!

I regret that the newer T.I. calculators chop the 13th digit off instead of rounding it off (by adding 5 and then chopping). Had the arithmetic and elementary functions been rounded in this way they would have been both more accurate and more nearly free from anomalies, more nearly like the hypothetical SR-52-III of §4.

Unclassified

Security Classification

DOCUMENT CONTROL DATA - R & D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) Computer Science Division Electronics Research Laboratory University of California at Berkeley 94720		2a. REPORT SECURITY CLASSIFICATION Unclassified	
		2b. GROUP	
3. REPORT TITLE AND NOW FOR SOMETHING COMPLETELY DIFFERENT: THE TEXAS INSTRUMENTS SR-52			
4. DESCRIPTIVE NOTES (Type of report and inclusive dates)			
5. AUTHOR(S) (First name, middle initial, last name) W. Kahan			
6. REPORT DATE April 6, 1977		7a. TOTAL NO. OF PAGES 22	7b. NO. OF REFS 0
8a. CONTRACT OR GRANT NO. ONR-N00014-76-C-0013		9a. ORIGINATOR'S REPORT NUMBER(S) ERL-M77/23	
b. PROJECT NO.		9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)	
c.			
d.			
10. DISTRIBUTION STATEMENT Approved for public release; distribution unlimited.			
11. SUPPLEMENTARY NOTES		12. SPONSORING MILITARY ACTIVITY Mathematics Branch Office of Naval Research Washington, D.C. 20360	
13. ABSTRACT This report responds to repeated requests for explanations of the arithmetic paradoxes perpetrated by the T.I. SR-52.			

Unclassified

Security Classification

Unclassified

Security Classification

14

KEY WORDS

LINK A

LINK B

LINK C

ROLE

WT

ROLE

WT

ROLE

WT