

## INTRODUCTION

p1.¶1: Since this CD-ROM isn't a murder mystery, we can tell you who-did-it right off the bat: *Nobody*. ...

- 0 Who invented the personal computer? / I did of course! Everybody knows that! Well, maybe not everybody; perhaps just those people who don't hold the prevailing opinion and are older than 107.
- 1 Wow! Superb writing throughout the entire work, now off to a flying start! It's great fun to read and to re-live those exciting times!
- 2 It's certainly true that no single person created from scratch the powerful little apparatus we now identify as a personal computer and have all come to love (or sometimes hate, I suppose). What we now have does indeed represent a transient merging of at least seven development streams: Concept, Architecture, Manufacturing Technology, Man-machine Interaction, Application Software, Communication Networks, and Marketing.
- 3 Yet I have to admit (despite my magnificently unassuming modesty) that your opening question's dramatic answer, "*Nobody*," given the way you have framed the question, is more than a little irritating. Had your question been either "Who invented personal computing?" or "Who created the personal computer industry?," I'd have had no problem whatsoever. The former can untroublingly be answered "*Nobody*," since all of the earliest computers were, at some point or other in their existence, used interactively -- at Whirlwind and Lincoln, in spades! -- by individuals for their own purposes; that is, for 'personal computing' as distinct from the 'impersonal computing' of batch-processing. The latter question can be answered "Jobs & Wozniak" and your story is then about Apples; or else, once again, "*Nobody*," since we all recognize that new industries grow from the convergence of multiple development streams. But as you have it, your theme immediately devalues the precedential importance of the LINC, which, so far as I know, was *the* point of departure from the streams of Concept and Architecture in accounting for the origins of the real-time, interactive, graphical personal computer we know today.
- 4 So I might as well say this up front: I think that in going on to develop your theme you have devoted entirely too much space to Time Sharing *per se*, which, after all, is just what using one's own computer is *not* all about. Yes, Time Sharing made it possible to use in personal computing mode the big machines that IBM had demanded be used only in the dehumanizing batch-processing mode. But using a computer 'personally' is not the same thing as using a 'personal' computer!
- 5 Stan, I've never claimed invention of the personal computer, nor did the IEEE on my behalf. But I was, so far as I know, the first person to *push* the idea of a real-time, interactive, graphical computer that didn't have to be shared, timewise or otherwise, to be viable, and then went on to put his money (the gov't's, that is) where his mouth was by designing and bringing into existence a sound example --- and, at that, an example having all of the structural features found in today's PC, however primitive the technology of those earlier days may now seem in retrospective comparison. [Such a sentence!]
- 6 I notice that Bell Labs doesn't get a mention here. I'll assume that the seminal UNIX activity, in which a great many of the now-familiar word processing tricks were developed quite early on, isn't to be part of the story because it's in the Software stream (though please note that most of today's PC-ers, who live in a DOS or MacSomething world, don't know of these origins). Even so, since you devote so much text to the Time Sharing epoch, this lack of referential inclusion of the very important UNIX work is unfortunate.

- p1.¶2: But this version of events, while as true as true can be, isn't what most people, even those in the computer industry, believe. ...
- p2.¶2: What is this tale but a contemporary version of David versus Goliath. ...
- p2.¶3: But did they invent the PC? Well, that's another matter altogether, because inventing means originating. ...
- p2.¶4: It's impossible to answer the question without defining it: What's a PC? ...
- p2.¶5: Obviously, there's no simple answer. ...
- p3.¶2: Our image, and thus our definition, of the PC has changed over time. ...
- p3.¶4: Clark was most certainly not a self-educated garage-bound tinkerer. ...
- 7 In what follows I'll try to behave myself and comment on your theme as given.
- 8 Without knowing your definition of "PC," I've no reason to be sure that your version is "as true as true can be" and not the unintended origination of yet another myth. There are indeed too many myths and re-inventions; but then, as Ivan recently reminded me, success has many parents.
- 9 With a few notable exceptions -- such as the magnetic-core memory, without which computers would have remained poor, stunted things until IC-memory came along -- we just did not, in those years at MIT and Lincoln Lab, patent things to the extent now practiced so enthusiastically.
- 10 In fact, I, for one, did feel in the early sixties that I was up against both the MIT and IBM Goliaths. [See my "steam roller" comment in the ACM's *A History of Personal Workstations* -- let me call it AHPW -- p358. For that matter, do try to contact Severo Ornstein for more on this. Tel: (415) 851-4258. He lives fairly near you at 2200 Bear Gulch Rd, Woodside. Lively, outgoing guy. Among other things, he founded the organization known as Computer Professionals for Social Responsibility. You'll like him a lot.]
- 11 Surely you could say something more like: "Some of the first PCs *in the form we know today.*"
- 12 Well now, so 'inventing' = 'originating.' Hmmmm...
- 13 I note that the modern PC looks about as much like the LINC as the Boeing-767, say, looks like the unquestionably-invented Wright brothers biplane. (I know, I know... powered flight was a *gigantic* step in transportation technology.)
- 14 Stan, if you are really intent on writing about the PC (a name probably owned by IBM, by the way), then maybe this is a good place to be more specific about your own definition. I urge you to think this through rather than sidestep the issue.
- 15 I am not now, nor have I ever been, an electronics engineer; 'a would-be physicist turned computerologist' is more like the correct descriptor here.
- 16 [Historical accuracy dept note, not actionable: The initial design and prototyping of the LINC were funded by Lincoln Laboratory. NIH didn't get into the act until 1963.]
- 17 I had no flair for electronics whatsoever. Computer-architectural design, yes. And thanks for the accolade, but actually there were so few computer designers in those days that we were all leading.
- 18 I believe Bells Labs built the first transistor computer. It did run for a while, though not very well; had the wrong sort of transistors.

p3.¶4: The scientists who used his machine seemed to love it; it was as personal as a computer could be back then. ...

p4.¶1: By today's standards, however, the Computer Society's decision seems incredible. ...

p4.¶3: Definitions change, especially when a technology is

19 My role at Lincoln, as I "helped to create" the TX-0, TX-2, and other transistor computers, was at the concept and architectural design level; the circuits and memory were the work of others, most notably Ken Olsen and Bill Papian. If any one person can be said to have built these Lincoln machines, it is Olsen, who was in charge of the engineering design and fabrication of all but the memory units. But I thought 'em up and produced the architectural designs. (My small staff gets the credit for much of the detailed logic design of the TX-2, and for some of its features as well). If you feel the need to be accurate without complicating matters, then perhaps you could adopt throughout something to the effect that I was the *architect* of the TX computers, Ken the *engineer* (as he had been of the MTC as well).

20 Designing the LINC was a considerably greater deal than designing the TX-0 was.

21 All right, so you couldn't carry it in your pocket like a comb. But as I've asserted above (note 5), it did have all of the structural features of the modern PC, though this is probably not the place to spell them out.

22 The handsome bronze medal from the IEEE merely proclaims 'First Personal Computer'; nothing about inventing it. 'Computer Pioneer,' sez the award.

23 Umbrage! Umbrage! Check this out: *By today's standards, however, awarding the Wright brothers an Airplane Pioneer medal seems incredible.*

24 Stan, please omit the reference to where I live (and work -- sorry I couldn't show you my office and its breathtaking, picture-window view of the city).

25 I object to "...and there doesn't seem to be anything personal about it." The Wright brothers' contraption did get us off the ground, after all.

26 The 1963 LINC's assembled at MIT cost about \$32,000 each, so said the accountants. (My target had been \$25,000.) I think \$45,000 may well have been a DEC price for either their assembled or kit-form LINC's; dunno about the price of the LINC-8 or its successor, the PDP-12, the form in which most of the world's thousand-plus LINC's were finally made.

27 All computers, unfortunately, still require a great deal of training to operate, though it is true that the first LINC's had precious few programs and documents to facilitate the task.

28 Hey! If the Application (ugh!) Software stream is to be part of your story, then your definition of PC must include: "...and has word-processing programs, spreadsheet programs, 'windows,' etc.," right?

29 Now I must rise to my full umbrageous height! Forgive me this parody of your final sentence: *If the Wright brothers' contraption were an airplane, then we might as well call a bicycle a compact car.* You couldn't even get a decent stewardess-served meal in the damned thing.

30 Seems to me that good definition is not impossible. And sure, we'll all look back on the '90s PC with amusement; but isn't your story about how we got this far?

undergoing a revolution. ...

p4.¶3: In short, the term—and thus the thing itself—is impossible to define, except in the vaguest way. ...

So what's *your* '90s PC? Enough of this vagueness!

I FIRST WIND

p5.¶1: History is a chain of developments, made up of an endless series of links. ...

31 And enough of my carping and complaining. From here on I'll try to be more constructive, vent less spleen. Can't promise, but I'll try.

p7.¶3: The vast majority of today's computers, including all personal computers, don't operate this way, of course. They're real-time, interactive, graphical machines ...

32 Careful! There is considerable disagreement over just when the modern computer era began. The folks in England won't much like your "definite starting point." Why not just say you're arbitrarily taking as your starting point the one described.

33 Ah so! Almost your definition of personal computer. Add a bit and you're there, though the terms 'real-time,' 'interactive,' and 'graphical' might still cry out for greater specificity.

p7.¶4: This machine was called Whirlwind, ...

34 Two more clues to your definition, I see: First, you imply that small size is a necessary characteristic of a PC. But this has two problems, seems to me: a) small with respect to what?; and b) if one were rich enough he might well have a Cray-3 as his personal computer. Second, you imply that a machine isn't a personal computer if it takes a cadre of experts to operate and service it. Well, I agree that it shouldn't take more than oneself to operate a PC; but until PCs are throw-away items, virtually all users must still depend on others to service them.

p13.¶2: It took about two years to design Whirlwind, and another three to build it. ...

p14.¶1: ...

35 Your figures for Whirlwind's multiplication process are correct. But I have never heard the term 'basic software' used in connection with this venerable machine and don't know where you got the figure of 35,000 lines of code. Its earliest programs were indeed written in octal notation, an equivalent of 0's & 1's and only slightly less onerous to deal with. But by about 1953, programs were written in an assembly language that dealt with words like 'add' and 'sub' and permitted simple if abbreviated alphabetic names for whatever variables were required in the task at hand; furthermore, simple compilers were just around the corner. Naturally, programmers were quite happy to give up 0's & 1's.

p14.¶2: (Incidentally, most of the computers of the period ... used words that were at least forty bits long ...)

36 Is this true? If I ever knew this fact(?) I've forgotten it.

37 What *isn't* true is that most of Whirlwind's interactive, real-time descendants also used sixteen-bit words. True, its immediate descendant, MTC, was a 16-bit machine. But its next descendants were the 32-bit SAGE machines (which could, however, also process 16-bit words), the 18-bit TX-0 and 36-bit TX-2 (which could also process 9-, 18-, and 27-bit words), and the 12-bit LINC. The DEC PDP-11 did have a 16-bit memory, but it used 48-bit instructions to operate on a variety of dataword sizes. Most of DEC's other machines followed my architectural lead and were based on 12-, 18-, or 36-bit word processing. Quite

independently, Seymour Cray also used the 6-bit modulus in his early machines.

## II SECOND WIND

p16.¶4: After awhile, Forrester realized that the scheme was sound in principal ...

38 It's a real pity the story doesn't include more about MTC, the Memory Test Computer. This marvelous machine, quickly but carefully built by a team led by Ken Olsen, was a simpler and vastly more compact version of Whirlwind. As its first serious programmer, and for a period of some weeks just about its only programmer, I tended to consider it my machine and personally computed a lot of good things with it. The unprecedented reliability and speed of its magnetic-core memory made using it a pleasure to one who had cut his programming molars on Whirlwind.

39 Where'd you get the "16K bank of cores"? True, 2 banks of core-memory, with 32x32x16 cores (~16,000) per bank, were hooked up to Whirlwind in '53, the first taken from MTC, the second enthusiastically assembled a few months later. But 2 banks of this size, in modern terms, comes to only 4 Kbytes of storage.

p25.¶2: Much to everyone's surprise, software turned out to be the biggest problem. ...

40 None of this software was written in 0's & 1's. The march up the language hill to higher levels had been going on since the mid-fifties, and good compilers were already in use.

## III THE REST OF US

p31.¶2: In 1957, IBM loaned one of its mainframes to MIT. ...

41 I believe the term 'mainframe' didn't appear until several years later.

42 Sorry I can't confirm Corby's estimate of the cost of a 704, though it sounds about right. Why not ask IBM?

p25.¶4: It's important to bear in mind that, in comparison to modern computers, you couldn't do much with the 704, the LGP-30, or any of the commercially available computers of the 1950s. ...

43 An arresting sentence! Damn, I wish my rare copy of the *1959 Weik Survey* hadn't been lost by a lawyer in Texas some years ago. It gives a remarkable overview of just where we stood at the end of the decade. It summarized the characteristics and production quantities of all the known computers, both academic and commercial -- and with pictures yet! It's an amazing realization that within the span of the first big decade we went from dumb punched-card machines to enough commercially available computers to warrant using the phrase.

p37.¶2: "The response time of the MIT Computation Center to a performance request presently varies from 3 hours to 36 hours depending on the state of the machine ..."

44 The term 'time sharing' came from engineering jargon at Whirlwind, where it appeared in a more restricted context: micro-time switching of different signals into the same computer element so rapidly that within the next-larger time span the element exhibited multiple personalities. The element was then said to be 'time shared' by whatever other elements produced its input signals or used its output signals. Not a very good term, but there it is.

p37.¶3: McCarthy wanted to modify the computer so that it could accommodate many people at the same time ...

45 McCarthy's version of time sharing -- I've always called it Time Sharing to distinguish the two meanings -- had to make its "points" in terms of the number of users that could be crammed into a shared-computer environment: *the more users the better*. Note that Time Sharing's efficiency thus depended critically

p37.¶4: In principle, time sharing, as McCarthy called the scheme, is a simple process. ...

on the extent to which the time being shared would otherwise go unused. Whenever this time was even momentarily in short supply, the users (McCarthy's thinkers/typers/try'er-outers) would, for that moment, be put into a mode of *competing* rather than sharing, and those tiny moments do add up. A more accurate name would therefore always have been Unused-Time Sharing. One of the most fundamental of Time Sharing's problems is that in any computer kept efficiently busy by numerous, eager users, very few of the computer's major elements have even a micro-moment to donate to unused time. The "points" formula increasingly becomes instead: *the more users the worse*, i.e., the slower and less powerful the resource -- down to less than that of batch operation, in fact, because the coordination of users' tasks takes time and memory-space -- though improved access is still a big gain. Corby would have a hellish job bringing conflicting expectations into reasonable balance.

46 In my view, the Time Sharing impetus, successful though it was in reasserting the value and legitimacy of personal computing, actually had the unfortunate effect of setting back the development of the modern personal computer by at least a decade. It's very disappointing, therefore, to see so much of your story devoted to it without some introductory observation that its manner of transforming a batch-processing computer into a real-time, interactive one would yield only very limited interactivity on a scale of time in which the reality was primarily that of the typing process. This limitation would be relieved only by putting more computing power -- not remote and shared but, instead, immediately proximal and dedicated -- into the hands of the user. That, inexorably, is exactly how things went in the ensuing years; teletypes gave way to 'glass teletypes' which, through the addition of more and more processing and memory electronics, grew increasingly powerful and eventually became efficient 'terminal/display processors' -- at which time it finally began to be noticed that these enriched terminals were doing nearly all of the work.

p39.¶4: To question this received truth made one highly suspect. ...

47 The small interactive machine that Jack Gilmore worked with was the TX-0. I've attached a copy of one of the pages from a 1959 Lincoln Memorandum that shows some of his work in progress.

p41.¶2: Things moved slowly back then. ...

48 Glitch? You have, "...to convert a time-sharing computer into a batch-processing one..."

49 Was McCarthy in a position to approve? Maybe just "endorse" would be more accurate.

p41.¶4: The hardest part of the job was the software. ...

50 In Time Sharing, stopping on a dime and switching from task to task was something that had to be programmed -- very slow switching indeed, with a dime as big as a parking lot.

51 I had already worked out a very efficient hardware solution to the context-switching problem some years earlier. It was to have been one of the main features of what I named the TX-1 (a machine that was never built). I was very pleased with my discovery of the architectural concept involved, which I considered a logical generalization of one of the ways multiple computers can be harnessed together, and later I did incorporate the feature into the TX-2 after we had built the TX-0 as a sound preliminary. I see that in the 1954 Lincoln Lab Memorandum describing my discovery, to which I had given the name 'multi-sequencing,' I wrote: *A multi-sequence program can also be constructed*

for a single computer. The general requirement is that the operation of one sequence must not interfere with the operation of any other. In general, this means that the operating registers of the computer must be time-shared by the sequences. I had already rejected the idea of the competitive mode in favor of the cooperative. These days, multi-sequencing goes (I think) by the name multi-programming -- or is it multi-tasking or multi-threading? I can never remember which. Incidentally, the multiple sequence principle was rediscovered nearly twenty years later by Chuck Thacker of Xerox PARC, who used it in the design of the Alto computer [AHPW p340].

#### IV SHARING TIME

p44.¶2: At Lincoln he was also involved in several small but important computer projects. ...

52 Maybe just a wee bit too cryptic? I guess the sense is Time for Sharing, right? (I've always stumbled over getting "The Secret Sharer" into the sense intended by Conrad.)

53 Now Stan, the program of the Advanced Development Group may have been small by divisional and overall Lincoln standards, but it was probably one of the largest R&D efforts then to be found in the computer biz. The TX-2 project was certainly not a small one in any case.

54 The MTC was built under Digital Computer Laboratory's auspices before the lab became part of Lincoln, though it was later moved out to Lexington as a Lincoln Lab property and continued in use there for several years.

55 Is this all there is to be on the light-pen episode? Sigh. I was very proud of this little gadget, which I had -- what would you say, invented? originated? made happen? -- specifically for real-time, interactive graphics.

56 The only people who don't have to know what they're doing are those who don't care about the consequences of what they will probably do wrong. (Remember, you heard it here first!)

p44.¶4: Starting small, the company's first products weren't computers but rather their building blocks ...

p45.¶1: ...

57 I presume you mean, "not knockoffs from the TX computer's *circuit modules*." But even so, you may want to adjust the thought communicated here. Gordon Bell writes about DEC's first products, the Laboratory and System modules, in these words: *The circuits used in both module series were based on the M.I.T. Lincoln Laboratory TX-2 computer circuits ... . All of the TX-2 basic circuits were used, except those gates which used emitter followers.* [Bell et. al., *Computer Engineering: A DEC View of Hardware System Design*, 1978 -- let me call it DECView -- p104]. I don't know whether or not the transistors used in these modules differed from those in the TX-2; but the mechanical packaging of the new DEC modules was certainly different from the TX-2's, and Ken's fine sense of style is clear in all three module forms.

58 "\$1.3 in revenues"? Incidentally, to help get things rolling, Bill Papian and I immediately bought a big group of Ken's very first building blocks, which he'd brought over to show us in salesman mode. I used them to mock up the control structure needed to operate the prototype LINC tape unit.

p45.¶3: In 1959, Olsen hired Ben Gurley, a brilliant and exacting engineer ...

59 ... and also a very warm, open, gentle guy, whose tragic murder only four years later grieves us all to this day. Among other things, Ben had designed the CRT display units for TX-0 and TX-2, and had taken over Ken's position at Lincoln. [I hadn't been able persuade him to stay with the group a while longer, even after a

long, last-try, hilarious evening we spent together in one or another of Boston's night-spots, in one of the darker of which he had indignantly responded to a demand that we produce our driver's licenses as proof of adequate age by pointing out to the waitress that between the two of us we had 10 children!]

<sup>60</sup> The PDP-1, though it had a more general input-output system and several new instructions, was essentially a production version of the augmented TX-0 -- the form in which the TX-0 had been sent to the MIT campus, and which had made it look very much like a transistorized MTC, which in turn looked very much like the basic Whirlwind. Ben had never designed a computer before and had little familiarity with how programming requirements affect design choices. He and I met at my hilltop home in Lexington that summer to work out some of the design details together.

p46.¶1: "The PDP-1 is one of the most important machines in the history of computers ..."

<sup>61</sup> Perhaps you don't fully realize that the TX-2 was many times the physical size, complexity, and performance of the TX-0. The TX-0 had taken me only a few days to think up and design in detail (architecturally, remember -- gates and registers, etc.) because it was such a primitively simple machine. With the tools available these days it would take only a few hours. The TX-2, on the other hand, was a multi-year effort and involved the detailed logic design work of a small group. Yet even if you properly limit your second sentence here by referring only to the TX-0, the miniaturization imagery is still incorrect, since the PDP-1 and the TX-0 were of just about the same physical size; moreover, the PDP-1 had a bigger CRT.

p47.¶3: "The PDP-1 was the world's first commercial interactive computer. ..."

<sup>62</sup> In going on to say that the PDP-1 "was the foundation for everything that followed," Fredkin fails to note that the immediate foundation for the PDP-1 was the MTC'd TX-0. Ah well. Does something have to be commercially available to be a foundation?

p48.¶4: The second PDP-1 was bought by ITEK ...

<sup>63</sup> I believe that Jack Gilmore was involved somehow, but don't remember the details. You can reach him to get the story by calling MITRE, where someone is maintaining a list of the Whirlwind old-timers. Charlie Adams never worked at Lincoln. He had been one of the principal programming gurus at Whirlwind back in its earliest days when Jack was employed there as a computer operator. (Yes, even Whirlwind spent some of its brief youth in batch-processing mode.) Some time later, Jack joined my TX-0 group for a very productive period of work. Later yet -- after Jack left Lincoln, I believe -- the two of them constituted the core of a small consulting group called Adams Associates, which may well have done some ITEK work under contract.

p49.¶1: ...

<sup>64</sup> Earl Pughe is also the man who supervised the modification and transfer of the TX-0 from Lincoln to the MIT campus.

p49.¶2: At about the same time as EDM was under development, a graduate student at MIT by the name of Ivan Sutherland was building a similar system, called Sketchpad. ...

<sup>65</sup> Lincoln did not donate TX-2 to MIT. The TX-2 remained in place until its retirement. Qualified use by MIT faculty and students who were prepared to come out to Lexington to use the machine was encouraged. Both Ken and I had been given MIT appointments as Lecturers in EE, the academic imprimatur that enabled us to supervise students. Ivan's work was done at Lincoln, where he -- like Charles Molnar and many other graduate students -- used the TX-2 for research as a student-associate of the Advanced Development Group. [Ivan has an engaging story to tell about how he timorously asked my permission to



p50.¶1: ...

undertake the Sketchpad work, and would surely impart it to you if you call him at SUN Microsystems Laboratories, Inc. (which, incidentally, is now directed by his brother William Sutherland -- Bert to his friends -- who also used the TX-2 in his doctoral research).]

<sup>66</sup> Ivan would also be able to tell you more clearly how EDM and the TX-2 Sketchpad compared. You're wrong about the lack of a display buffer in Sketchpad. The displays flickered only when a picture that had lots of figures in it was being displaced or magnified, for the reason that the TX-2 wasn't fast enough to program big updates to the buffer without stuttering, but were otherwise pretty stable. This effect can still be seen in some of the PC application (ugh!) software packages available today.

<sup>67</sup> Ivan co-established the E&S Corporation with Dave Evans.

p50.¶3: Despite the success of Corby's little time-sharing project, for example, some people at MIT had serious doubts about the usefulness of time sharing. ...

<sup>68</sup> I don't know who the other serious doubters were, but here's *my* argument: The claim was that Time Sharing would efficiently put the power of a big computer into the hands of Everyman, and do it inexpensively at that. Unfortunately, the reality could only be otherwise. Consider that throughout the time interval or intervals during which any Time-Shared or otherwise-shared computer works on a single user's task, that user must effectively pay for the use of whatever computer resources the execution of his task prevents *others* from using and paying for, since there ain't no free lunch. This means that he must pay for all of the available resources that his task *doesn't* require along with those that it *does*. This is why, for example, running a small problem on a large computer generally incurs unnecessary expense, though the cost often goes unnoticed or is charged to convenience.

<sup>69</sup> Now consider a job that actually requires, but is unable to fund exclusively, most or all of the available resources of some computer big enough to Time Share. In Time-Sharing execution it must either a) unacceptably tie up the machine during specially-assigned large intervals, or b) take an unacceptably long time to run when competing on more even turf with numerous other jobs. What this meant was that MIT's Time-Sharing jobs, on average, had to be relatively small if very many users were simultaneously to be given access without too much queuing delay. Fine for McCarthy's thinker-typers, not so fine for his tryer-outers when their programs needed a lot of speed and memory to run without ho-hum; interactive graphics, of course, was out of the question.

<sup>70</sup> The net result was that the big Time-Shared computer -- already made slower and less memory-capacious by its co-resident operating system -- could offer to its simultaneous users, in relation to their numbers, only what amounted to lesser, expensive "machines" of relatively poor efficiency. Yes, there were important economies of scale, notably in memory size, that helped to restore lost points, and the operating system did include a large, useful, ever-growing suite of new compilers and other programs that alone justified the size of the machine. Nevertheless, the mismatch of expectation to performance is manifest.

<sup>71</sup> No doubt Corby realized all this. And yet, note the trace of the dream in MAC = Machine Aided Cognition when the reality was MAC = Multiple Access Computing.

p50.¶4: Within MIT, a few people, most notably Wesley A.

<sup>72</sup> I infer from the draft's typo ("helped designed") that your first go at the introductory sentence here omitted the word 'helped.' [See note 19 above.]

Clark, who helped designed the TX computers, didn't care for either time sharing or batch processing. ...

p51.¶1: ...

p51.¶2: In April of 1961, the Long Range Computation Study Group, as the faculty committee was called, submitted its report ...

p51.¶3: For Teager not only went along with the recommendation that MIT get a mainframe and convert it to time-sharing, but went a big step further ...

p52.¶1: ...

p52.¶2: While Teager was dreaming of PCs, Clark, Charles Molnar, and a few colleagues at Lincoln were actually making one. ...

p52.¶3: Prof. Arturo Rosenblueth, a portly man who ...

73 I doubt that there was anyone else at MIT, among those who were aware of the issue, who didn't believe that Time Sharing was going to be the greatest thing since sliced bread; but then memory plays tricks on the biased. Who ya got? In any case, there was no 'They' in your "They wanted MIT to look into the feasibility of personal computing." Nor did I want MIT to look into anything; I simply proceeded from conviction and had the resources to ignore the steam roller and get on with it. [Severo would certainly have things to say about this (see note 10).]

74 I first heard the term 'personal computing' from Alan Kay at PARC, where I spent a week or so every month or so during its formative years.

75 The term 'mainframe' had not yet appeared, though it certainly would have been understood: a big-deal computer that is to be considered far more important than any lesser machine nearby or subservient.

76 I'm glad to learn that Herb -- hadn't he actually chaired the committee? -- did write a dissenting report. Apparently I should have written one as well. He and I were, I believe, the only members of the committee who refused to sign its final report.

77 The phrase *personal computer* (as you emphasize it) in Teager's dissenting report may well be the first appearance of the term. When was this written? It would be fun to see it sometime and compare his sketch with the ones in my own 1961 LINC design notebook, which includes a great many badly-drawn pictures.

78 By going along with the recommendation while visualizing his gadget as a remote input-output console, Teager missed the mark in one important way: a console subservient to a Time-Shared machine is not a personal computer. Well, not in my book, anyway, though Teager certainly had the right objection in mind. Had his advice been taken, years might have been shaved off the ensuing epoch of enslavement to the Big Machine. Presumably it was Teager's conviction that resulted in just such a console at MIT a couple of years later. I believe it was built by John Ward and was called the Kluge or maybe Kludge, but I don't know whether it worked out well or not. [Interestingly enough, Jerry Cox and I, in teaching a class in computer design at Washington University in 1965, had the students build a small gadget like Teager's but also operable in stand-alone mode. We called it the 'PC' (!) for Programmed Console. Jerry later refined the design for more emphasis on stand-alone operation and made a bunch of these little machines, which were then placed around the country in a sponsored evaluation program very much like the LINC's.]

79 Teager got it almost right. A personal computer has *no* usage formalities to contend with.

80 NO! Stan, re-read the book! *Prof. Walter Rosenblith!* And maybe he's portly now, but a better description of the man I knew in the '50s and '60s would be, "slight, energetic, cosmopolitan, etc." He was also a cautious man, and the history was not as you have it here. He never approached me for such a thing

as a LINC, and, in fact, had rejected the idea of having the TX-0 moved to his lab at MIT. Nor was I ever assigned to look into it; the LINC, like the earlier Average Response Computer (which I *had* been asked to make), was permissively bootlegged within the larger program of Lincoln's Advanced Development Group. And it wasn't that Walter 'wanted' anything like a LINC, but, rather, that I knew what he and his group and many others like them *ought* to want, and certainly needed. [See AHPW p357.]

81 There's that "helped designed" typo again.

p53.¶1: With financing from the National Institutes of Health (NIH) and other government agencies ...

82 [See notes 16 & 80 above concerning who paid the bills when.]

83 Just a few adjustments for precision, if you care: 1) of the four console boxes, one played only the role of the empty card-slots found in today's PCs and was often put somewhere else; 2) *alphanumeric* keyboard rather than *typewriter* keyboard -- and it was used for typing in not only commands but also text and symbolic data; 3) display *CRT* rather than display *oscilloscopes* (we tried two CRTs for awhile but gave up one for the final version); and 4) it isn't that the LINC 'had' a tall cabinet, but rather that its electronics and power supplies, being inconveniently bulky, were held in a tall cabinet that could be put out of the way and connected by cables to the console parts you worked with -- just as we often put the equivalent cabinet of some of today's PCs out of the way and cable it to the parts we work with.

84 Important point: the magnetic tape reels anticipated today's floppy disks in a way you haven't mentioned (yes, I know, you can't say *everything*). The tapes were pre-formatted into fixed cells and blocks, just as we now do with disks. This critically important departure from all other tape-design practice of the day was something I had worked out for the TX-2.

p53.¶2: In May 1962, Clark demonstrated the prototype at the NIH in Washington, D.C. ...

85 The NIH guy was deadly serious. Even though he could see the cat attached to one end of a cable and the LINC to the other end, he seemed to feel that the setup was a hoax since there was no intermediating punched-card machinery -- the *sine qua non* of his computer world.

p53.¶3: Although the LINC had a meager 1K of internal magnetic-core memory (expandable to 2K), ...

86 1K *12-bit words*, that is, not the 1-Kbytes now generally associated with "K."

87 Is the LINC not really a personal computer because it was designed with real-time laboratory work in mind? Or because its memory was so small?

p54.¶1: ...

88 Graphical interaction was provided by 8 knobs, the interpretation of whose settings was under program control. The knobs were conveniently located on the CRT unit. The LINC, like all of its MIT predecessors, also had a useful audio output.

89 Where'd you get the 4000 hours? (I don't dispute this, just curious.) Incidentally, the last(?) functioning LINC classic was retired from service just this year!

90 The tape drives were really very special and were the most significant technical innovation in the LINC. Nothing like them had ever appeared before, except for the TX-2's giant, less-well conceived version. I credit their ultimate reliability entirely to Charlie Molnar's fine engineering.

p54.¶2: With the exception of the tape drives and the console, the LINC wasn't, as Clark acknowledged, technically exceptional. ...

- 91 Yes, I once did demonstrate the ashtray trick to NIH visitors. Even the prototype design was so conservative that the unit would still work with a piece of thin paper inserted between the tape and the read/write head, as I demonstrated as well.
- 92 The person who suggested the spilled-blood spec was kidding, but he'd made his point.

93 The console *module* (box), that is. I designed this module so that the user could interrupt the machine smoothly at any point in its calculations to see in detail what was going on, then proceed step by step until his understanding was improved enough to resume operation at full speed. I felt that this kind of facility was necessary because of the many uncharted waters in the real-time control of experiments. The console provided other functions as well, and designing for smoothness was extremely difficult. All of the early MIT machines had consoles of one kind or another, replete with switches and indicator lights for coded display of the various machine states; but the feature of miscue-free smoothness was uniquely the LINC's. Consoles have disappeared from modern machines, which only incompletely compensate for the loss (to those who need this kind of detailed understanding in their work) by programming means.

p54.¶3: In comparison to today's PCs, the LINC doesn't seem the least bit personal. ...

- 94 Aha! To be a personal computer it has to be small and easy to use? If you feel this to be an essential part of your definition, add it in. But please note that no early computer of this kind was either as small or as easy to use as its successors would be.
- 95 You ask, "Why wasn't the PDP-1 a PC?" My simplest answer would be that its design wasn't based on the "single-user-as-master" philosophy, as William Calvin put it in a 1982 letter to BYTE magazine (though the implication evident in the attached copy -- that the LINC was the first computer one could use personally -- is clearly wrong). A personal computer, like a personal pocket-comb or a personal anything-else, is a computer you don't have to share with anyone else unless you want to. The PDP-1 could indeed have been owned in this way; but, so far as I know, it wasn't.
- 96 Of course, other elements must be added to my simplistic definition to bring it into consonance with our understanding of what constitutes the modern PC. Gordon Bell, DEC's foremost architectural designer in the days of the LINC and still one of the most authoritative in the computer biz, ventures a richer definition [AHPW p9]: *A personal computer or pc is a self-contained computer with secondary file memory and appropriate transducers to interface with people.* The PDP-1 would therefore not be a personal computer, since it wasn't, in Gordon's terms, self-contained. But then in going on, he permits himself the loose equivalence, 'personal use of a computer = using a personal computer,' modern informality's too-facile indistinction that merely contrasts personal computing with the nasty old, impersonal, batch-processing; thus his definition continues, *A personal computer is used interactively by one person at a time [my emphasis], at a location convenient to the user, and may "belong" either to the person or to a group.* This permits him to assert later (p30): *The PDP-1 continued in the tradition of the "MIT personal computers."*; and yet in further explaining what he considers a personal computer to be -- and almost requiring "LINC-like" to be part of his definition -- he writes the following: *The microprocessor, memory, and mass-storage technology appearing in 1975 lead directly to the personal computer industry [my emphasis]. Early computers utilized the simple, single*

process, stand-alone operating systems developed for both interactive, time-shared computers and stand-alone minicomputers. Nevertheless, the first personal computer, the LINC [,] was built in 1962, long before its predicted technological time. If Gordon had believed the 1959 PDP-1 truly to be a personal computer, I've no doubt he would have said so.

97 The official DECView is silent on the PDP-1's status. Its only reference to 'personal computer' is this: *The [LINC's] tape system and a powerful CRT-based console made possible the first complete personal computer available to a user, in this case the researcher, at a reasonable price.* [DECView p175]

98 Note, incidentally, that only 50 PDP-1s were made (cf note 101 below). [DECView p166]

p54.¶4: The NIH set up a program to distribute the LINC throughout the biomedical community. ...

99 Actually, the biomedical *research* community.

100 Not *Rosenblueth's lab* but a specially constituted lab set up by MIT with Lincoln's help, and with *Rosenblith* and William Papian serving as co-directors. And the visiting scientists assembled their LINC's from a kit of almost-ready-to-go parts (just a bit of corrective re-wiring had been necessary). The only disassembly required before shipping home amounted to disconnecting the cables.

101 I am not now, nor have I ever been, an employee of the Digital Equipment Corporation. Your closing shrift here is very short indeed, and your numbers need adjustment. True, about 50 "classics" were assembled, 21 of them by DEC. But there were also 143 LINC-8s, the form I worked out (as a consultant to DEC) with Richard Clayton, with the intention of sustaining DEC's interest a bit longer by wedding the LINC's logic to the PDP-8 then compelling the company's greater interest; and DEC then made 1000 PDP-12s, Clayton's re-do of the LINC-8 [DECView p176]. In addition, a good many Micro-LINC's, which used early ICs, were made by the Spear Corporation. [DEC went on to put DECTape -- its nearly-exact copy of the LINC's tape unit -- into many of its machines, and made them in production quantities.]

p56.¶1: A community, perhaps even a culture, grew up around CTSS. ...

102 Yeah, the dream was indeed hampered by the reality. A friend of mine defines a personal computer as a computer that doesn't run any faster at night!

p56.¶2: Many practices that are now commonplace on computer systems were started by CTSS, ...

103 ... but not the use of both upper-case and lower-case letters. Early versions of computer-adapted IBM Selectrics had already appeared; one was used on one of the early LINC's.

## V LICKLIDER'S DREAM

p61.¶4: Lick got an LGP-30, which was one of the first, and most popular, small computers on the market. ...

104 "(The LGP-30 is described in detail in Chapter II.)" -- ??

p62.¶2: For Lick, using the PDP-1 was akin to "a religious conversion," ...

105 [Perhaps all that Coke was part of the religious experience. Lick developed a serious caffeine addiction to it, which he had to try to correct from time to time.]

p63.¶4: Since computers had a long way to go before they could do these things, ...

106 Flexowriters were not used in the vast majority of computer systems as you imply, and even the systems that did make use of them did so mainly for off-line, punched-tape preparation.

p64.¶1: ...

p70.¶4: Work in man-machine interaction was already underway at various places on campus ...

107 Please don't neglect to mention the TX-0 in your final sentence here. The main reason Papian and I sent the TX-0 to the MIT campus was that we wanted to put a real-time, interactive, graphical computer into the hands of students and faculty. It was heavily used.

p71.¶1: ...

p71.¶3: Even though it was part of the Pentagon, DARPA was a small and casual place, and could act fast. ...

108 "\$2,220,000"? Nuthin! Within the next few months Rosenblith and Papian had managed to secure a \$27,000,000 grant from NIH for extension of the specially-constituted lab's work (see note 100 above), with LINC-like activity as one of its initial centerpieces. [See also AHPW p370.] What happened to it, and how MIT managed to drop the ball, is another story for someone to tell some day.

p76.¶2: In the meantime, Bell Labs dropped out of the effort ...

109 Did you know that the name UNIX is a reactive word-play on MULTICS?

(To be continued ...)

110 Wow again! Stan, you're telling a great, absorbing story, and I'm truly impressed by the scope of the undertaking and by how beautifully your writing keeps the historical momentum growing and developing. I sincerely hope it will receive overwhelmingly more enthusiastic praise from the misinformed masses than it does the inevitable righteous objections of the few. But I also hope that you will have squarely faced the thematic issues I've raised by the time your work sees the light of CD-ROMdom.

111 Finally, I hope that by now you'll have forgiven my frequent exasperation and didacticism, at least enough to want to show me the chapters that finish the story.