HOW TO GO OUT OF A BUSINESS AND LEAVE THE CUSTOMER HAPPY
By Bruce Perkin

Since the 1960's, Librascope had a distinguished career of making disk memories. Not the type of disk memory in common use today; the disk memories that Librascope promoted were of the fixed head variety. Having one magnetic read/write head per track gave them faster average access time than the more conventional disks - such as seen today in your own computer - with a single head on a moving arm. This was because the access time was only that required for the disk to revolve until the desired data was beneath the read/write head - no more than one revolution of the disk. There was no additional time delay while a single head would be repositioned to a new track. Furthermore, with many heads reading or writing in parallel, the data transfer rate was much higher. Librascope manufactured all shapes and sizes of fixed head disks, from giant units containing six disks of four foot diameter providing 12 recording surfaces and using about one thousand heads, to units as small as 10 inches in diameter. See Figure 1 below.

Figure 1. A Librascope Press Kit from the early 1960's era. Note the 4 foot diameter disk.

Bear in mind that nothing made that long ago can compare with the performance of the modern moving head disk that takes up a mere corner of your laptop computer for capacity, access time, or data rate. This story is about how the Company got out of a dying business yet continued to support the product.

When, in early 1986, Carl Sorensen became responsible for the product line, the disk memories were of more modest size that fit within a standard relay rack enclosure and were used in Navy shipboard systems that handled highly classified information. In that type of application the disk memories performed well because of their good data reliability and the fact that the stored data was non-volatile so that the computer system it was embedded in could be shut down or re-booted and return to normal operation once the programs and data had been re-loaded from the disk to the computer. The fast access offered by the head per track design was also useful because it permitted rapid swapping of data between the disk and the computer's internal memory - very important for analyzing large data sets that exceed the size of the computer's internal memory.

The downside to being a head per track disk memory manufacturer was that the manufacturing process was difficult. A whole array of mostly hand assembled read/write heads was required. Each of the 16 parallel bit streams required its own read/write electronics and those had to be switched 16 at a time to an array of 256 read/write heads by selection electronics. This made the electronics complex. Finally, the disks and their plated magnetic coating had to be nearly perfect for all of the tracks to store data reliably. The disk plating was outsourced, reducing the control the company had on the product. A few spare heads and tracks were provided but these were insufficient for very large disk defects or for very many small ones.

Sorensen shortly hired Karl Fetterhoff. Karl had been previously exposed to the commercial world surrounding the recent phenomenon of personal computers and was very knowledgeable about the state of the art and was one of the early computer hobbyists. Observing the disk memory problems, he came to realize that the capacity of the disks could now be matched by the use of solid state dynamic random access memory (DRAM), something that was totally infeasible back when the disk memories, then in production, were designed. If that substitution were made, he reasoned, we would need neither more of the difficult mechanical parts nor any more critical mechanical assembly.

One day, in early 1987, Don Barton – who managed the disk programs, Fetterhoff, and Sorensen were meeting when the concept came together ... Barton was complaining about the manufacturing problems, Fetterhoff volunteered that he thought the rotating memory could be replaced by a solid-state design, and then Sorensen said "Hey, if that were possible, we could submit a VECP*!!! The rest became history.

*VECP = Value Engineering Change Proposal

Based on its premise, the VECP "RAMdisk" had to have all of the mechanical and electrical characteristics of the original disk. That meant that the same aluminum housing needed to be used, and it had to attach both to the shock mounts that held it in the cabinet and the cabling that connected it to the controller. The motor didn't have to whir – that wasn't in the specification. Nor was there any specification about what the unit looked like inside, so we were free to do as we pleased there. The obvious conclusion was to use the aluminum outer castings as they were, modifying them only as necessary to secure the new parts inside, and hope that they all fit inside the housing in the end. The original purpose of the rigid aluminum castings was to protect a rather delicate disk memory from shock, vibration, and dust. The electronics used to simulate the disk memory were much more rugged. The cast housing can be seen in Figure 2.

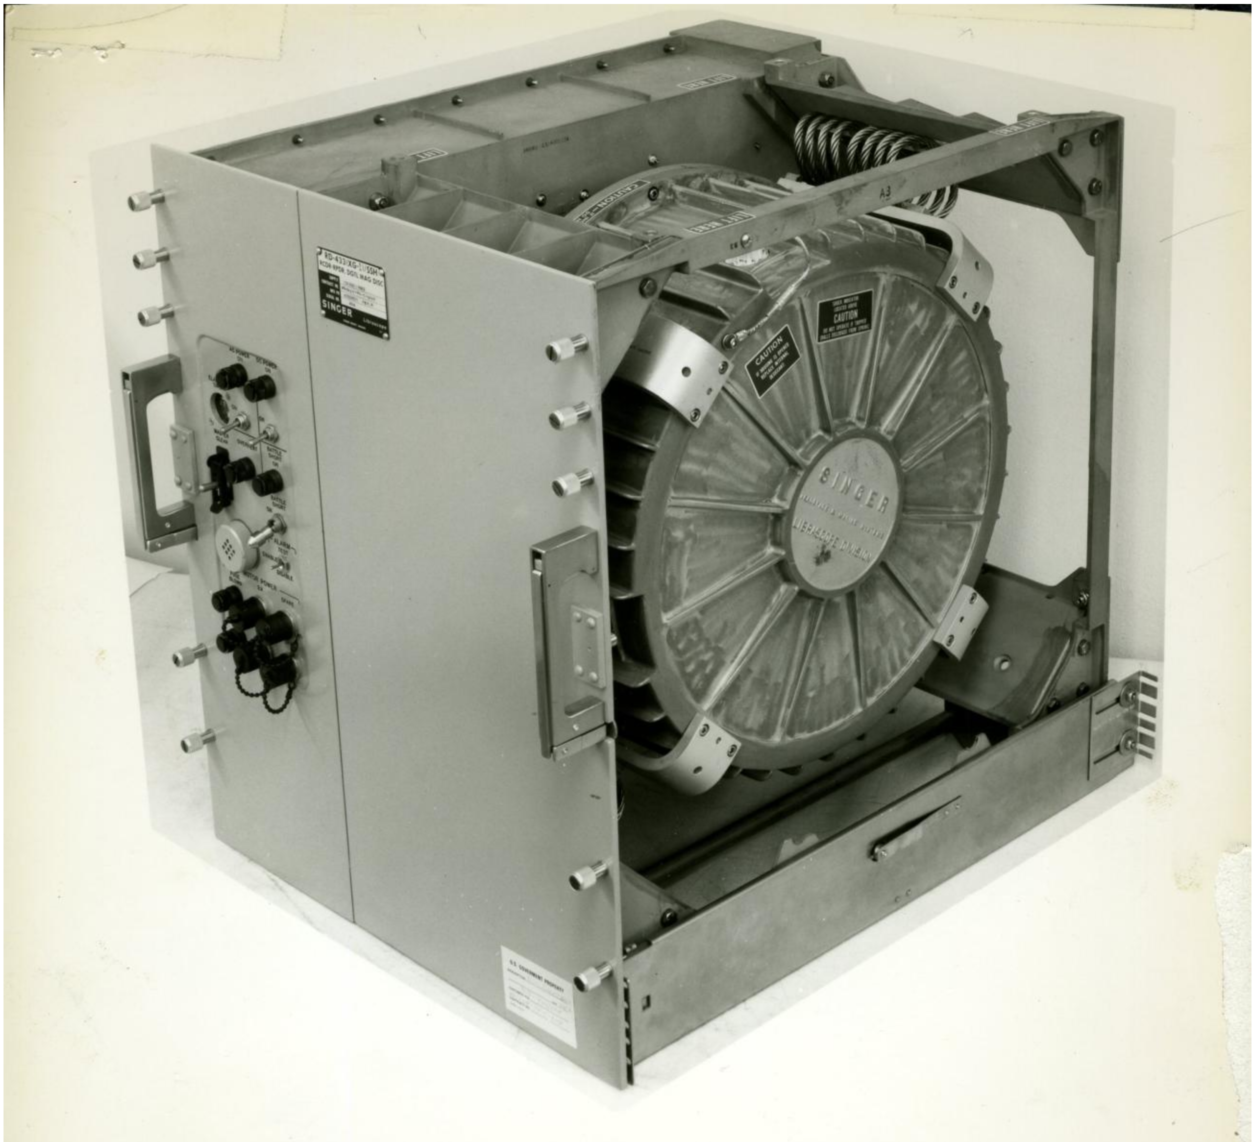

Figure 2. RD433 Disk Memory. Note, it is not possible to determine, from the photograph, whether the casting contains a rotating disk or a RAMdisk. The motor (RAMdisk storage battery?) housing, on the far side of the casting, is not visible.

Fetterhoff soon went to work to overcome objections to his design. Without power, the data stored in DRAM would be instantly lost – the opposite of the desirable disk characteristic of non-volatility. Karl found that in the Navy, power loss was relatively rare. In fact, the vessel is in serious trouble when power to its systems fails for any extended period of time. The only time power might be off would be when the system containing the disk was down for repair. That time, he was assured, would be short. So, he added a storage battery to power the DRAM array when line power was off to his design concept along with the charging system it needed, and estimated there would still be enough space within the disk housing.

The interface required by the DRAM was different from that of the disk. He was sure that this wouldn't be difficult for Howard Stahle and his logic designers. Howard assured him that there would be no problem making the new DRAM device look to the computer exactly like a revolving disk; in fact, the computer could operate faster because the access time could be made equal to the shortest expected disk delay and would have none of the extra delay that formerly occurred while the disk revolved to where the data was beneath the heads.

The third objection was that the Navy was adamant that all Integrated circuits had to be Mil-Spec and had to be packaged in ceramic rather than in plastic. The problem with the use of ceramics was that the clays that were mined to make them had occasional atoms of radioactive isotopes mixed in that could radiate onto the memory device that it encapsulated. The alpha particles from those atoms hitting the DRAM IC surface could flip a memory bit from a one to a zero, making that bit of data erroneous. No permanent device failure was done, just a bit flip. The problem was much greater for ceramic than for plastic IC packaging. Magnetic disk memories were immune to this effect. But a DRAM IC memory could lose one bit of data and make a large data field invalid or make a program not work at all. This was serious. The data on disk could be easily checked and restored dockside, but once underway, the data had to retain its integrity for the entire mission at sea which could be several months. Fetterhoff showed that, for a memory of this size, that there would be a single bit upset, on average, every 500 days or so and concluded that there was no need to worry for typical missions at sea.

I became involved at that point because it was an engineering go. The first thing I noticed was that 400 to 500 hours of average failure free time did not mean that we were assured mission reliability. In fact my computations showed that for a mission reliability of 99.5% that the mission had better not last more than a couple of weeks, much shorter than the typical mission at sea.

On the other hand, if we were able to correct, on readout, for two errors in a data sector, the average time to perceived data failure would be expanded almost indefinitely. It took a small amount of added DRAM to store redundancy data to make error detection and correction work. But it was a radical change to the logic used to make the DRAM simulate a disk. John Gustafson worked out the redundancy encoding and Howard Stahle's logic design crew made the logic for error detection and correction using, then, state of the art field programmable logic arrays.

Everything still managed to fit within the odd shape of the original mechanical disk. To the mechanical and circuit board designers, it was easy compared to the original mechanical assembly. The VECP was approved by the Navy for procurement.

To my knowledge, no problems ever occurred with the Value Engineering change. It was applied to all of the contracts that Librascope had for disk memories making it possible to convert the entire disk product line to RAMdisk. Because the RAMdisk was less expensive than the conventional disk memory, the Navy saved millions of dollars. Librascope's share of the savings amounted to over one million dollars. Librascope left the customer fully satisfied and supported, and never touched a rotating disk again. In recognition of the VECP's success, Sorensen arranged for all of Librascope's RAMdisk customers and all of the projects major contributors to receive a special commemorative plaque from the Company. Fetterhoff, the originator of the RAMdisk idea, went on to contribute greatly to the innovation of other programs.

Ironically, Sorensen later attended a Navy seminar, at which he was able to call attention to the success of the RAMDisk VECP, with the idea of seeing its expansion into other disk memory supported Navy systems. Vice Admiral Tuttle, who was the guest speaker at the seminar, invited Sorensen to meet with him at the Pentagon. The meeting went well. But, when the Government's own value engineering staff got hold of the idea, they decided to go further and eliminate the requirement that the replacement had to have all of the form, fit, and function of a disk memory assembly. Instead, it would be no more than a self-powered card rack. Librascope never got any of that business.