VIEWPOINT

Football and Microprocessors

How the Microprocessor Industry Can Benefit from Football's Experience

By Nick Tredennick

So what could football and microprocessors have in common? For a few years, they shared a common problem, but since there are about 10^n football fans (where *n* is the number of people interested in microprocessors) the problem got solved first for football. The *n* of us interested in the microprocessor industry could learn a lesson from the football industry. Here's the story:

If you are old enough, you probably remember Monday morning arguments about NFL quarterbacks:

> "Did you see what Baltimore did to Washington yesterday? Even on his worst day, Johnny Unitas could whip any other quarterback for quick-release, completions, and play calling."
>
> "No way, man. Did you go blind before the Dallas game? Pound for pound, Eddie LeBaron had no equals on the field."
>
> "You guys are both crazy. Didn't you see Minnesota make Green Bay look like a high school JV? No other quarterback has ever come close to Fran Tarkinton for scrambling ability and total passing yards."

And so on.

No more. The arguments are gone—and it isn't because all those football nuts got pink slips and went to work at Jiffy-Lube. The arguments died when the NFL Quarterback Rating System was born. The arguments have all been settled. All you have to do is look in the *Information Please Sports Almanac* for any given year and read who the top quarterback is. It's all based on a formula. The formula is based on completion percentage (C), yards per completion (G), number of touchdown passes (T), and number of interceptions (I). The formula probably looks something like this:

$$R = \alpha \times C + \beta \times G + \delta \times T + \sigma \times I$$

I don't know what the coefficients are, but it shouldn't be too difficult to figure out. No one bothers to figure it out because no one cares. You just look in the book to see how many positions Dan Marino moved up since last year's publication. No football fans are likely to sit around arguing about whether the parameters are appropriate or whether the coefficients are properly weighted.

That's the end of the story. It's how the NFL solved the question: "Who's the best quarterback?" The computer industry has a similar question: "Who's the inventor of the microprocessor?" It's been over twenty years since the first microprocessors were produced and we are still arguing about who invented the microprocessor. I suggest we learn a lesson from the success of the NFL and help resolve the controversy with the introduction of a "Microprocessor Inventor Rating System."

For the past few years I have hosted an awards ceremony at the annual Microprocessor Forum. In 1990, I collected a list of five candidates for an award to the inventor of the microprocessor. For a similar award in 1991, I collected a list of twelve names. In a single year, my list more than doubled. Even at this, I suspect the list is still in its infancy. With the dramatic improvement in East-West relations, I'm sure we will soon be hearing about twenty or so inventors of the microprocessor from the Eastern Bloc countries. Also, since most of the rest of the world's inventions can be traced to China or Egypt, I expect the imminent discovery of works by another twenty or thirty inventors. In fact, it wouldn't surprise me if some untranslated hieroglyphics turned out to be a CMOS cell library. I hope this trend isn't following Joy's law: $2^{(year-1984)}$. If it is, we will see the same heart-stopping annual increases in the size of this list that we have been observing in the performance of Sun workstations since 1984.

The list of candidates is large and shows every indication of continuing to increase. Could they all have made fundamental contributions to the invention of the microprocessor? Unlikely. I believe there is a theoretical upper bound on the number of inventors of anything. I haven't seen any fundamental research on the topic, but when DARPA finally funds the work, I'll bet the number turns out to be either *e* or π. Whatever it is, it probably isn't Avogadro's number—or as my spelling checker wanted to correct it, Avocado's number—which is where the list appears to be headed now.

But we don't need a theoretical solution, all we need is an engineering solution. I'm going to follow the current trend in computer science by skipping the theory and going directly to quantitative procedure. I think the Microprocessor Inventor Rating System can give us the engineering solution we need. Here's the list of parameters:

- Patents
- Patent Priority
- Publications
- Publication Citations
- Commercial Products
- Relevant Experience
- Professional Acceptance
- Self Promotion
- Aggressiveness of Public Relations Staff

The formula looks very much like the formula for the NFL Quarterback Rating System, with a term for each of the parameters multiplied by some arbitrarily selected coefficient.

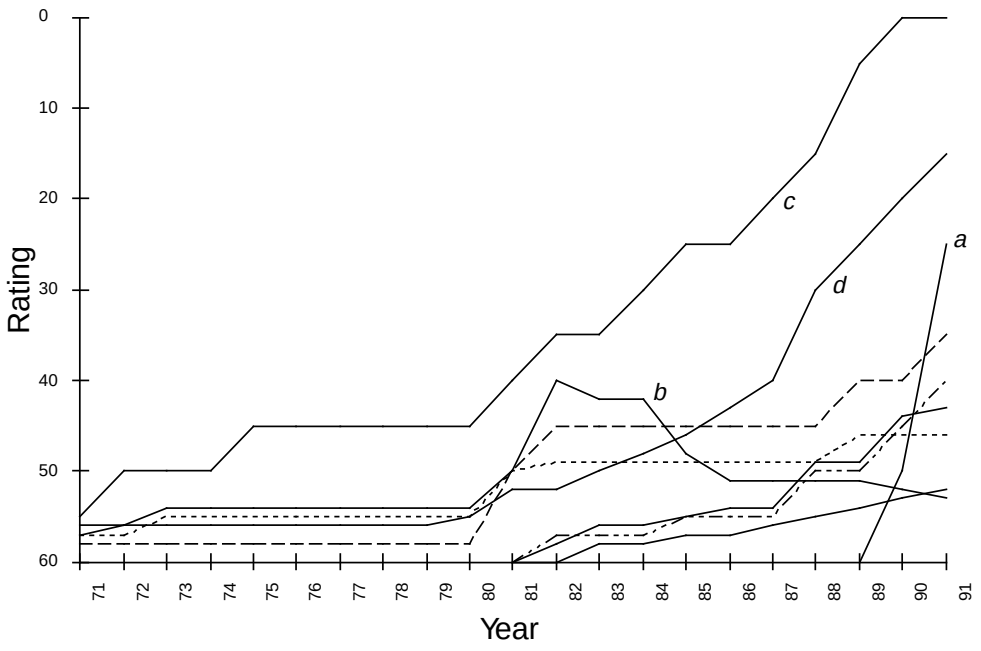

Figure 1. "Inventor of the Microprocessor" ratings from 1971 to 1991.

- ***Patents and patent priority*** means about what you would think: it's the number of fundamental patents on which a candidate's name appears as an inventor and the date the patent application was filed. This category always contains potential surprises, partly because the progress of applications in the patent office remains secret during processing. Some patent applications have spent twenty years in the patent office prior to the granting of a patent.
- ***Publications and publication citations*** includes technical works by the candidate as well as the number of times those works are cited by others discussing the invention of the microprocessor. Citations by authors of their own works are counted at a reduced rate. Publication citations also includes the number of times a candidate is cited as the inventor of the microprocessor in technical literature. Works in which the author claims credit for the invention of the microprocessor are excluded, since this is generally accounted for in the self promotion category. (Students of sensitivity analysis may see a weakness in the formula due to the potentially dramatic effect of the exclusion in certain individual cases.)
- ***Commercial products and relevant experience*** are important. If you were in the third grade in 1971, you probably didn't invent the microprocessor. If you worked in the machine tool industry in the early 70s, you probably won't get as much credit for your experience as someone who worked for a semiconductor company.
- ***Professional acceptance*** doesn't count for much, but it attempts to measure whether a particular inventor-candidate would be a believable choice among computer professionals. Medical advice from TV actors playing doctors (*e.g.*, Marcus Welby, Dr. Kildare) has always carried more weight than advice from the Surgeon-General—except among doctors.
- ***Self promotion and aggressiveness of public relations staff*** are largely a concession to reality. History has proven the best way to get credit for something is to be very skilled at taking credit for it.

These parameters might just as easily be used in a formula to determine the inventor of the blender or the modem, but we are concerned with determination of the inventor of the microprocessor and, therefore, have to consider the influence of dominant characteristics of the candidate pool in setting the coefficients of the various parameters. These parameters are obviously not independent. There is, for example, a close tie between self promotion and having a public relations staff and a possible connection between self promotion and publications.

Figure 1 is a plot of the ratings of a small number of candidates from 1971 through 1991. I have left the names off the chart because the results are preliminary—and to avoid embarrassment and litigation. This chart shows some of the difficulties facing someone attempting a quantitative solution to the "Who invented the microprocessor?" question. Perhaps the foremost difficulty is with the potentially unlimited number of candidates. The NFL, in rating quarterbacks, considers

only NFL quarterbacks with at least 1500 passing attempts—a closed system growing slowly with time. In contrast, there is no such barrier for entry to consideration for inventor of the microprocessor.

As you see from this chart, there wasn't much interest in who invented the microprocessor during the 70s. The ratings in the chart for the 70s mostly show the effect of working on integrated circuit design and publishing some of the early technical papers. The 57 engineers actually doing the work at ten companies probably thought it was obvious. You can see the increased interest beginning in the early 80s. The figure is incomplete, since I was able to collect and plot only a small fraction of the recently identified contenders, but it does show some interesting features. (Perhaps Miami Dolphin's quarterback Dan Marino used a similar plot of his ratings improvement and its projection in negotiating his recent $25 million five year contract.) For example, the line beginning in 1989 that appears to be rising faster than the number of instructions in the MIPS architecture, shows the effect the patent office and a large public relations staff might have on a contender's rating (see line *a* in Figure 1). This line may represent the only candidate on the list with the equivalent of the NFL's 1500 attempts requirement, except in this case all of them are correspondence with the patent office.

Line *b* on the chart, showing a preliminary peak in the early 80s and tapering off later, may indicate the possible ill effects of over-aggressive self promotion. Other lines show consistent long-term or belated efforts to improve ratings either by individuals or, in at least one case, by an aggressive legal department at a large company which may be embellishing the patent claims of one of its former employees (see lines *c* and *d*, for example).

Unfortunately, I don't have any final results yet. I'm still collecting data and tuning coefficients, but progress is slow because the project is unfunded and, therefore, must be relegated to "hobby" status. I suppose I should be filing grant proposals with DARPA or NSF to get funding to complete the project, but I'm not from a big name university and this doesn't look like the beginning of a fad, so its chance of being funded is nil. It's too bad, because I think this kind of work is important to computer science and society.

If the intense acrimony built up over the last fifty years in the competition for recognition as inventor of the computer is any indication of what is in store for the candidates aspiring to recognition as the inventor of the microprocessor, we owe it to society and the giant candidate pool to head off the coming street fight. We owe it to ourselves and our profession to establish rules for friendly competition and a definitive, quantitative result. And may the best candidate win! ♦

Embedded Processor Trends

Continued from page 13

established. Some very-high-end embedded processors will continue to simply collapse more of a general-purpose system onto a single chip. Others will be designed with a specific application area in mind; these will have a set of peripherals peculiar to their intended use. Still others will consist of older processor cores augmented with additional processing units that make them appropriate for modern applications and let them exploit the cost advantages of simpler implementations.

Increasingly dense IC processes combined with better automatic design tools will lead to less time spent "hand tweaking" logic and circuits. In turn, a new processor derivative can proceed from definition to implementation in less time. Automatic design tools might not give the optimal implementation, but they help get silicon functioning correctly in less time. Consequently, the pace of innovation in the embedded market can be expected to accelerate.

The ability of vendors to produce customized embedded processors will have to improve, given the fact that the number of embedded applications will also increase at an accelerating rate. In ten years, personal calculators went from being expensive tools for engineers and students to being disposable "freebies" available to anyone with a promotional need. In the future, small televisions, telephones, and even personal computers may be relegated to such disposable status.

The golden rule of embedded control is that low cost wins. The most successful members of the 960 and 29000 families are not those that have higher performance but those that have reduced costs. In the 29000 line, the newest family members show performance moving sideways while cost moves down.

Over the next several years, as fabrication technology permits the integration of several million transistors on a low-cost die, embedded processors will provide a tremendously rich field for innovation. Innovation in general-purpose processors will consist mainly of increasing cache sizes, more pipelining, instruction decode and execution units capable of processing multiple instructions per clock cycle, and possibly multiple processors on a single chip. Embedded innovation will include all of the above plus increasingly complex mixtures of processors and peripherals, support for fuzzy-logic or neural-network processing, perhaps some field-programmable logic, and ever-expanding on-chip RAM and ROM.

In short, while general-purpose processors must evolve in ways that are compatible with existing application software and long-lived operating systems, embedded processors will evolve with the products into which they are embedded. ♦