ELLIOTT PRECHTER EWAVES Flash

elliottwave INTERNATIONAL

April 2014 Issue
© April 16, 2014

# MUST GO FASTER

EWAVES needs to be fast. In fact, out of all of the ongoing improvements to the program, speed is near the top of the list—for several reasons.

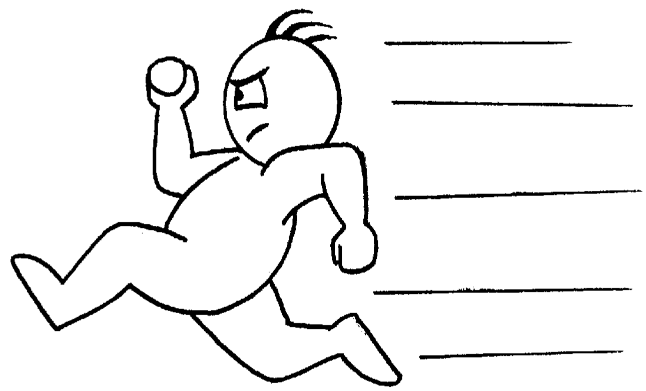

Above all, the faster EWAVES runs, the more detailed an analysis it can perform in the same amount of time, giving us higher quality wave counts. The better the wave counts, the higher the edge provided by each signal, which directly translates into better-performing recommendations.

A faster EWAVES would also allow us to introduce an intraday level of granularity into our computer-based analysis. Having the ability to finesse the wave structure at a very small degree will eventually allow us to provide Flash signals intraday. Currently, our *U.S. Intraday Stocks Pro Service* (http://www.elliottwave.com/wave/USIntraday) provides intraday analysis for those trading the S&P, Dow and Nasdaq.

Finally, speed is crucial for research. Past issues have noted that even with the latest computer processors, some of our more expansive backtests (covering multiple markets over long time periods) take *weeks* to perform. The most time-consuming tests are our mandatory *quality* backtests. Each time we modify the analysis engine, we must analyze a massive array of markets and calculate statistics in order to confirm the modification in fact improved results. We also compare the newly outputted analyses against a set of known wave counts to ensure that they have been maintained. If we can speed up EWAVES sufficiently to turn weeks of quality testing into mere days (or less), then the rate at which we improve EWAVES 1.X/2.0 will increase commensurately.

## The Challenge

To say that running EWAVES is a computationally expensive process is an understatement. It's a massive program.

Why should it be so difficult and take so long for a machine to provide analysis, when a human being can often just glance at a chart and label it within minutes? There are two answers. First, a human will never be as rigorous as EWAVES. It begins with no preconceptions, has no vested interest in a particular outcome, and is not influenced by questionable fundamental or technical indicators. It operates using only the purest form of Elliott wave analysis.

Second, computers do not operate like humans do, at least not yet. A computer is like a super-calculator, able to perform billions of operations per second in sequence, yet without understanding or intuition. The human brain sports nearly the opposite design. It is akin to a massively parallel processor with slow sequential throughput, running at just 200 Hertz (200 sequential operations per second). Some of the hardest problems for humans, such as factoring 10-digit numbers, are easy for modern computers. Likewise, some

P.O. Box 1618 • Gainesville, GA 30503 USA • 770-536-0309 • 800-336-1618 • FAX 770-536-2514

of the easiest tasks for people are still unattainable by machine. For example, PhDs specializing in computer vision are still a long way off from getting computers to recognize objects.[1] Meanwhile, if you hand a photo to kindergartners, they can immediately discern all the objects in it, and still have time to explain why the class's pet rabbit is secretly a super-hero.

Technology has made progress in emulating the human brain, yet it has far to go. Until pattern-recognition technology catches up, the best way to perform wave analysis by machine is to *formalize* the process in a step-by-step, calculator-style manner. In other words, we need an algorithm.

**Crunch-Time, Baby**

Elliott wave analysis is a *hard problem*. I don't mean "hard" in the sense of a difficult task to perform well (though that also is true). Rather, I am using the definition from complexity theory, in which a problem that has no known fast or efficient solution is considered computationally hard.[2]

EWAVES must perform extensive processing to produce a full Elliott wave analysis. To understand why this is so, consider the price chart below of the Dow from 1987-2003. Although only a few wave labels end up on the chart, the amount of detail that it must consider is large. And that detail increases exponentially as it examines smaller degrees.

This means that an in-depth analysis takes a really long time. So how can we make it faster?

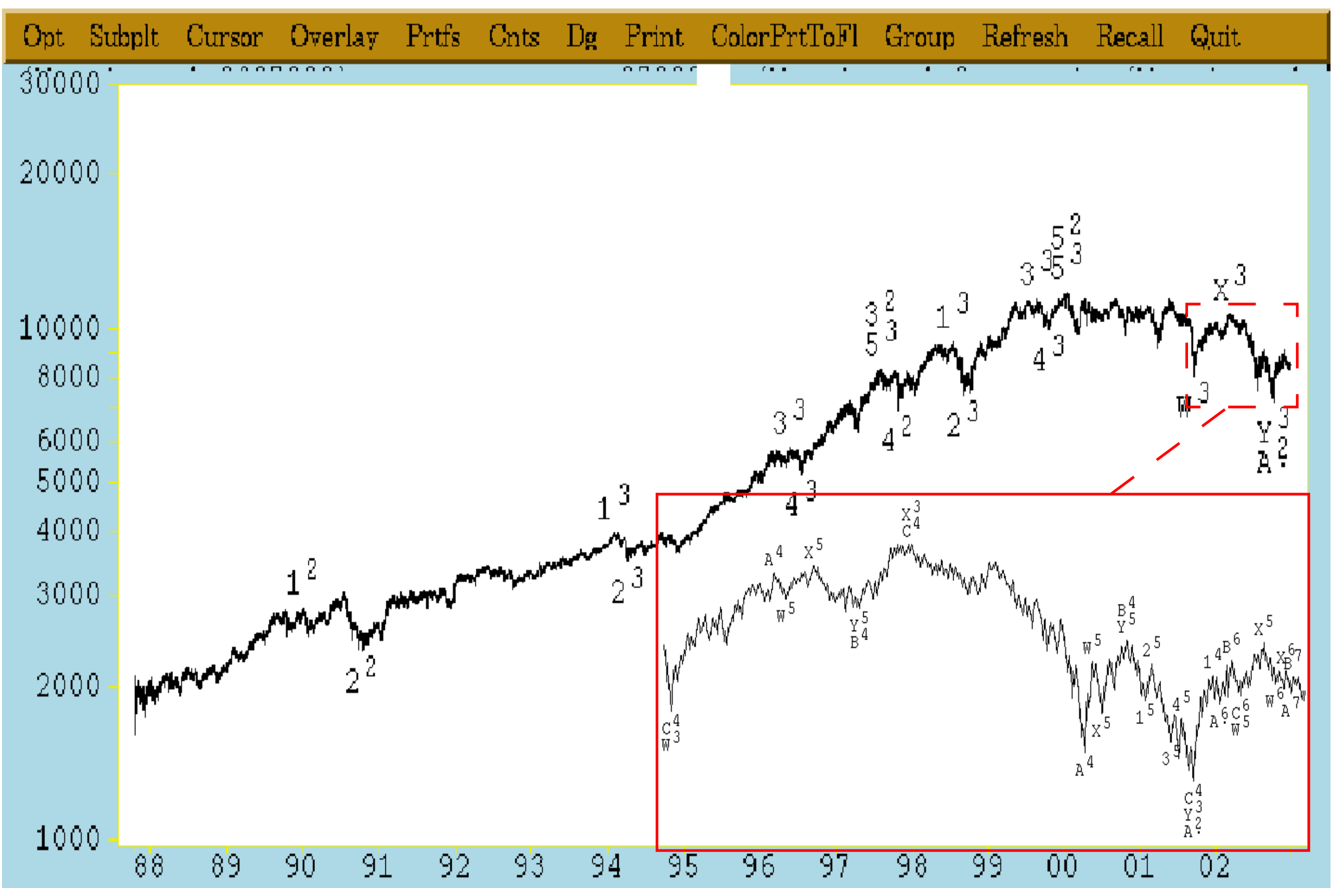


*Dow Jones 1987-2003*

**Work Harder**

One way to make EWAVES faster is to ensure that the computer uses all of its available resources to tackle the problem.

A decade ago, PCs typically had only a single central processing unit (CPU) core. A CPU core is like an assembly line: It receives instructions sequentially and executes them one at a time. These days PCs have multiple CPU cores, each one capable of receiving and processing a separate sequence of instructions. Each core still performs only one instruction at a time, yet multiple cores can work simultaneously.

In theory, each additional core should increase overall processing speed additively. In practice—alas—adding more CPU cores to a computer does not automatically provide improvement. Here's why.

Imagine that our CPU cores are cooks in a kitchen. In our example we have four of them, as in the image below. Let's say their job is to bake a delicious cake. (Mmmm, cake...) Cook one is in charge of reading the recipe, cook two prepares the dry ingredients, cook three prepares the wet ingredients, and cook four preheats the oven. Four cooks who divide the labor across multiple tasks will produce the cake sooner than one cook who has to do everything.

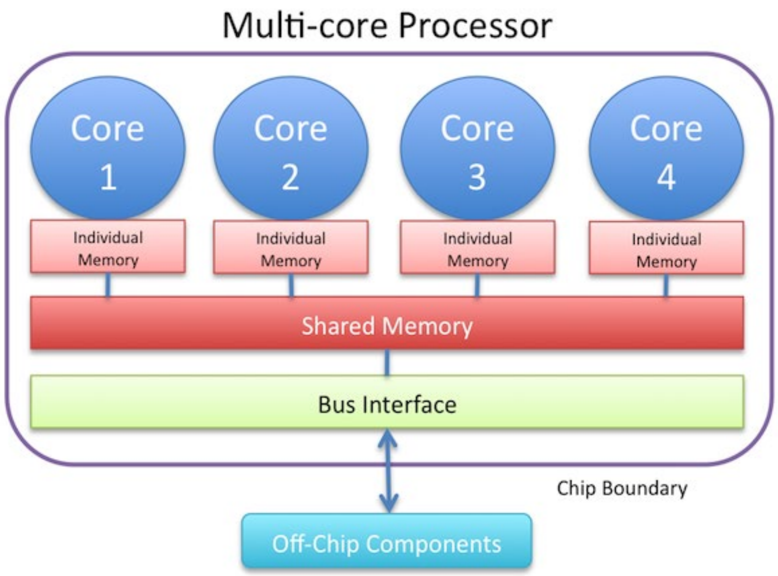


But no matter how many cooks are in the kitchen, they will not be able to mix the cake *and* bake it at the same time. The baking step is *dependent* on the batter mixing step. In other words, some operations must happen in a sequence and cannot be done simultaneously. So, there is a fixed number of cooks for maximum speed in baking a cake. Beyond that, more cooks will not mean faster production. If we replace the word "cake" with "wave analysis," then the issue we face becomes clear.

The January 2014 issue of *EWAVES Flash* explained how we have recently revitalized the EWAVES project. The original Lockheed code—written in the era before multi-core PCs—still comprises the bulk of the EWAVES 1 program. It operates mainly in a sequential, step-by-step process rather than doing simultaneous work. This design impedes making any one *particular* analysis run faster by throwing more cores at it. Fortunately, EWAVES 1 can still benefit somewhat from using multiple cores, because we can run *different* market analyses simultaneously on the same computer. But since we cover a limited number of markets, with just a few computers we hit a limit at which point no additional machines will allow us to produce analyses any faster.

The good news, however, is that EWAVES analysis is not really like baking a cake. The analytical work is (theoretically) highly divisible into sections that could be performed in unison. That type of framework would benefit from thousands of CPU cores working simultaneously.

That is exactly how we're designing EWAVES 2. It is being built from the ground up with the capacity to leverage an unlimited number of CPU cores for *each* individual market analysis. A currently standard 4-core CPU has the potential to increase the speed of each EWAVES 2 analysis by 4x. And going forward, it's likely that 8 or even 16-core CPUs will become standard soon. The more cores we have, the faster EWAVES 2 will run.*

**Assuming optimal cache-locality.*[3] *We are currently optimizing the usage of CPU cache in EWAVES 2.*

*Factory assembly line for Royal Wedding Chocolate Cake—Now that's some serious cake throughput!*

**Work Smarter**

"Work smarter, not harder" is a useful adage. In my experience, *how* you approach solving a problem can often save more time than trying to race through a routine. In other words, a tortoise will beat a hare if it knows the right shortcuts to take.

Consider, for example the parallel trend channels that bound impulse waves, illustrated at right. Once you know the ends of waves 2 and 4, then you can use the channel to project the end of wave 5.

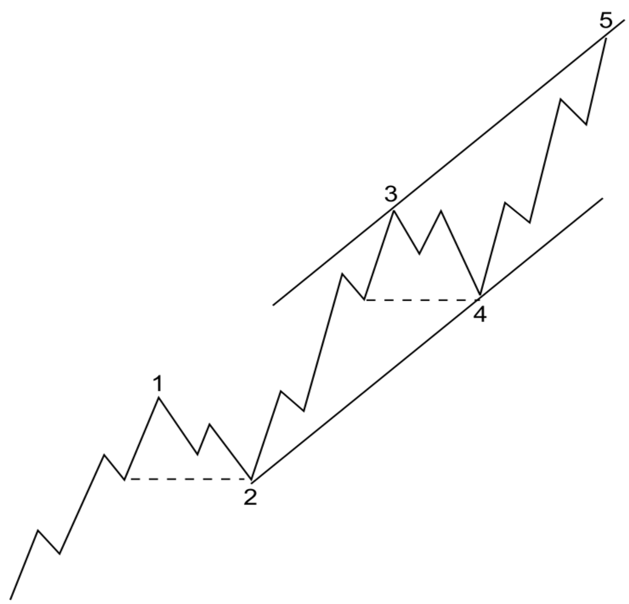


EWAVES "sees" trend channels (and similarly relevant constructions) when it calculates qualitative wave structure metrics. It usually does so by identifying where a particular line intersects with price. For example, when it projects the top channel line forwards, does it hit the top of wave 5? If not, when and at what price does it intersect, both before and after the top of wave 5?

A human analyst answers these questions by using a ruler. But a computer can't use physical tools. So to emulate a ruler in EWAVES 1, the Lockheed programmers did what most financial applications do: Test each price bar one at a time, using the line's slope to check for an intersection. Sounds simple.

But what if we're looking at a really large amount of data—say, 100 year's worth. In that scenario, it may take a while to find the intersection. In a worst-case scenario, the machine would need to check every single price bar to know if and where an intersection occurred.

This kind of one-by-one, sequential testing is okay for simple point-and-click charting programs; in fact, to an observer, the results usually appear instantaneous. But a program as sophisticated as EWAVES may require checking millions of line/price intersections over the course of an analysis. The time adds up. Our research shows that this process costs about an hour's time for each EWAVES 1 analysis. There has to be a better way.

**Space Partitioning**

It turns out we *have* come up with a better way. Take a moment to imagine just how much work goes into checking each price individually: It's like asking "Where is my toothbrush?" and getting a reply that says "Check the entire surface of the planet, one square inch at a time until you find it." Clearly, it would be best to ignore that advice and just check your bathroom drawer.

The reason most of us can find our toothbrush so quickly is because humans naturally think hierarchically when dealing with issues of space. Our toothbrush is in a certain drawer. That drawer is in a certain bathroom. That bathroom is on a floor, in our house, which is in our neighborhood, and so forth. By thinking hierarchically about locations in space, we're naturally performing a form of what is known as *space partitioning*.

Although we still have the entire surface of the world available for our search, our mental partitioning of the world allows us to instinctively ask high-level questions that immediately exclude most of the planet: Is the toothbrush in China? No, never been to China—so *bam*! We just eliminated millions of square miles of search space.

Space partitioning can be easily adapted to find the intersection of a line and geometry, as EWAVES requires. If you've ever enjoyed 3D virtual reality software, then you've unknowingly benefited from this technique. Have you ever walked into a virtual wall and wondered how the software instantly knew to stop you? The program doesn't check every wall to see if your path intersected it. Rather, the program had already divided the space in its virtual world hierarchically, as in the simple diagram shown below. Then it checks whether you may have hit something within a large region of space; if you couldn't have done so, it quickly rejects that region and all of its sub-regions. It does this at each level in the space hierarchy so it can quickly converge on the exact geometry that your path intersected. Instead of slowly testing every one of perhaps thousands of walls in the world, space partitioning allows the program to find the intersection almost instantaneously.

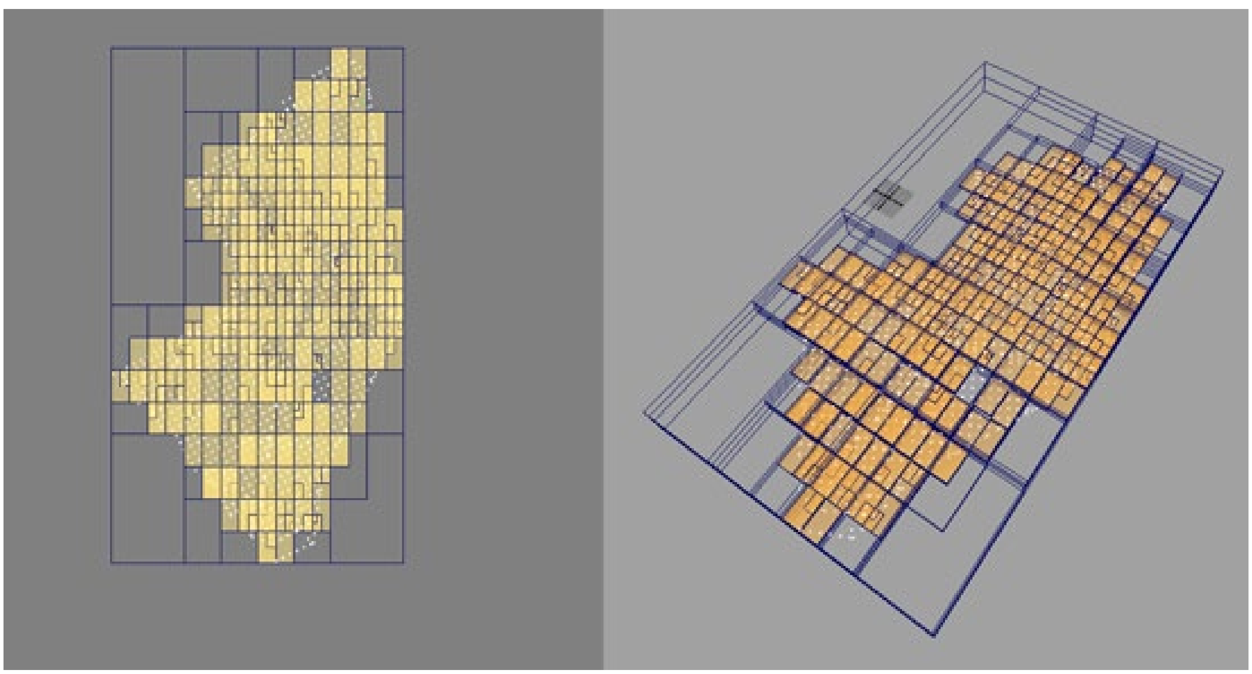

*Space Partition Hierarchy, Common in Real-Time 3D Graphics Engines*

To speed up the line/price intersection code, EWAVES 2 uses a space partitioning technique similar to those employed by real-time 3D graphics engines. (Without getting into specifics, we only need a 2D version.) Our alpha builds of EWAVES 2 show that space partitioning reduces the time required for certain aspects of an analysis to less than a cumulative second vs. the cumulative hour for the same task required by EWAVES 1. This equates to a speed increase of 3600x for certain sections of the program, an improvement equal to decades of hardware evolution. Although space partitioning is common in 3D graphics, we are fairly certain that we are the first to use this technique for financial analysis.

**Signal Review**

(by Greg Pyron)

In the last Signal Review, we looked at the initial stop loss order and how the EWAVES count determined that level. This time we focus on a more dynamic example of stop adjustment: how stops are currently modified in response to real-time changes in EWAVES' analysis.

As EWAVES updates the wave count each day, the most recent smaller-degree patterns determine the stop. Our recent long recommendation for crude oil provides a good example.

Figure 1 shows EWAVES' count for oil on January 30, 2014, when we issued the recommendation. The initial stop is at 91.23 (a tick under the wave $B^5$ low), a break of which would invalidate the top wave count. As sometimes happens, this initial stop is pretty far away. As the market moves, EWAVES provides us with additional wave structure, and in turn the strategy adjusts the stop.

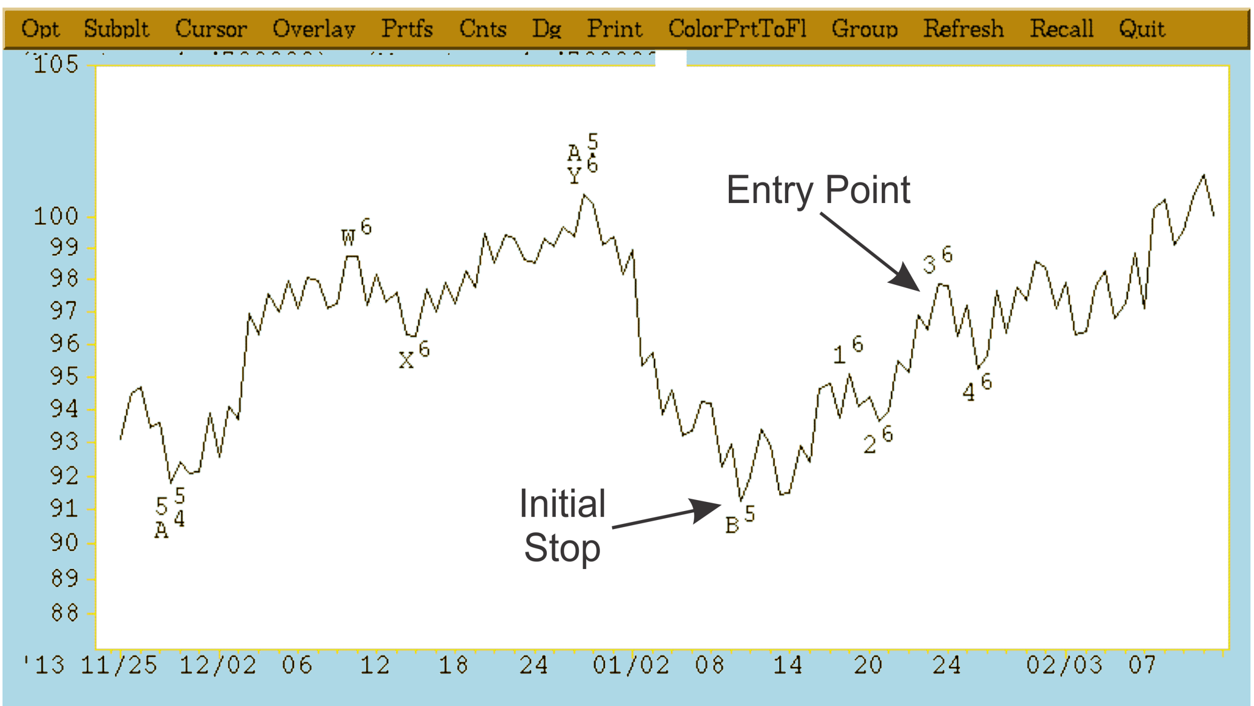


*Figure 1*

On February 17th (Figure 2), new information allowed the stop to be raised to 99.30. The reason is that this level intersects the trendchannel defining the wave.

Figure 3 shows EWAVES' count on crude two days later on February 19th. With more labels now in place, the stop was raised to 100.54, a break of which would have exceeded the expected retracement for the substructure of wave $5^7$. Finally, on February 24th, crude rallied to a new high for the move within what EWAVES was labeling as wave five within an extended fifth wave from the low back in January. Therefore, the stop was raised to 101.20 in anticipation of the rally's end.


*Figure 2*

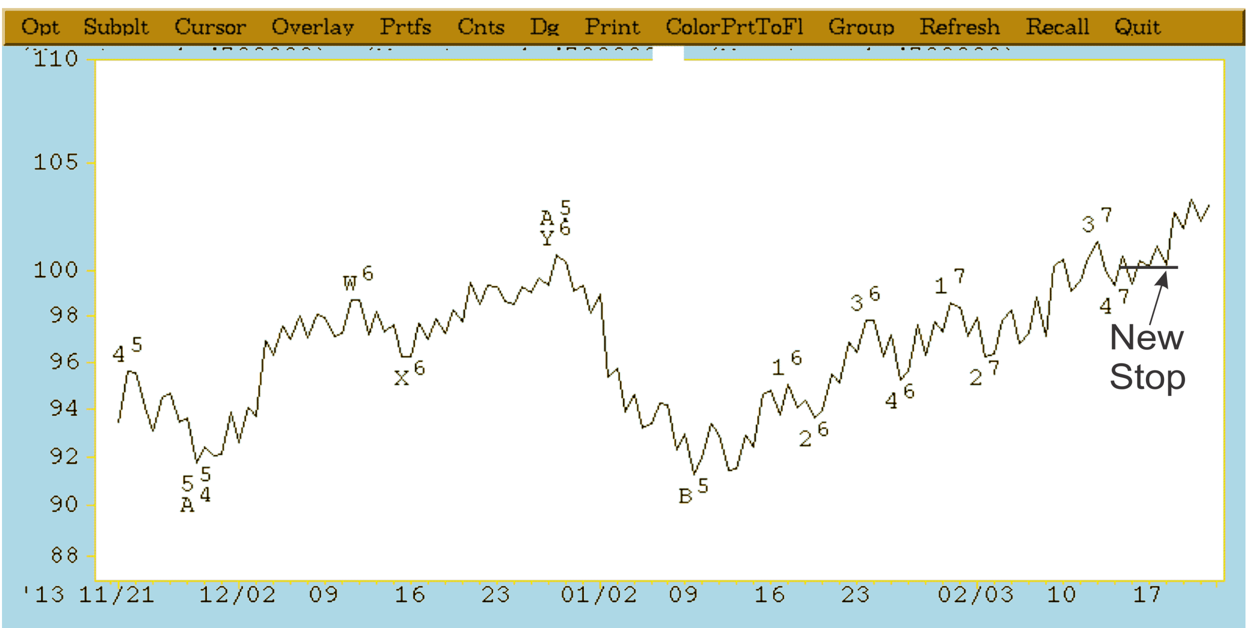


*Figure 3*

The position was stopped out the next day, locking in a four-point gain. Crude prices had one more surge but then reversed sharply, quickly retracing most of February's rally and underscoring the benefit of our adjustable stop placement.

**Iconic Counts**

(by Dave Allman)

A guess-and-check approach in EWAVES' development really comes into play when we make occasional Occasionally EWAVES produces an analysis that would not pass as a beauty pageant winner. Sometimes EWAVES generates a count that—to a human—just doesn't have the "right look."

Consider, for instance, a wave count where the fourth wave is disproportionate to the rest of the structure. One would like to think that a minor adjustment would generate the desired result. And one would often be wrong! Why? Because making just one change may create other problems, and each change will affect other markets. It's a complex issue. But the answer is simple: check, check and check.

Seemingly sensible changes may in some cases radically alter the wave count. And more often than we would like, sometimes a change prevents EWAVES from producing a reasonable count at all for a given market.

Enter the iconic counts. We have a catalog of unequivocally clear impulsive and corrective Elliott wave structures from about a dozen markets, which we call *iconic* because they are cast in stone. No subsequent market action can change how these markets are properly labeled according to the Elliott Wave Principle. Whenever we adjust the code, we check the new iteration of the counting engine by running it on our portfolio of iconic counts to ensure that they remain intact.

Let's look at one of these iconic counts. Here's a chart of cash gold from 1980-2000. The drop from \$850 to \$250 is corrective, period. Nothing gold does now or in the future—be it drop to zero or rise to infinity—will change this conclusion.

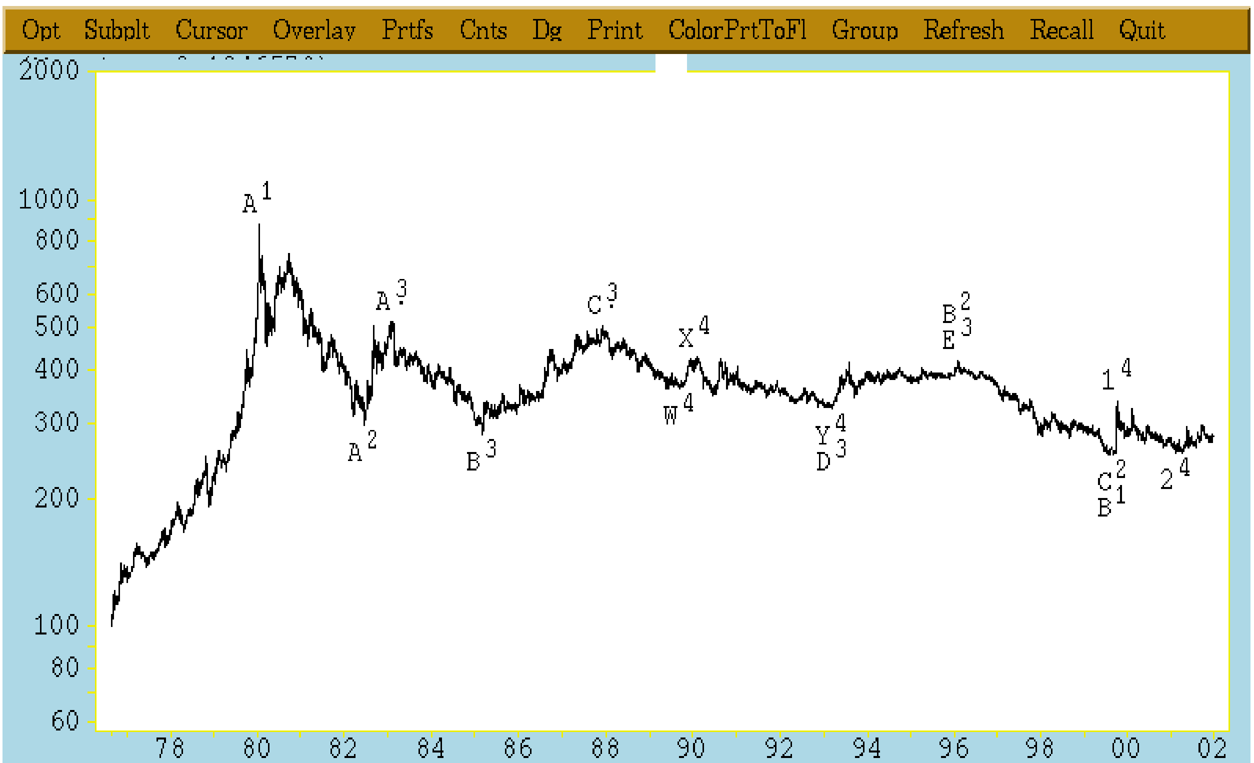


*Gold (Cash prices) 1980-2000*

The decline is a simple zigzag, labeled $A^2$, $B^2$, $C^2$. But look more closely at wave $B^2$—a contracting triangle, roughly spanning the years 1982-1996. Having lived through that market, I can tell you it was sometimes a bear to count (no pun intended). Yet I remember like it was yesterday when Bob surmised a triangle in the wave B position—a simple and elegant solution to a bunch of back-to-back-to-back, three-wave structures. You can imagine our excitement when EWAVES—after crunching the data at multiple degrees for hours and hours—subsequently labeled the move a triangle!

Suppose we were to tweak the counting engine order to improve EWAVES' overall counting ability but after checking the suite of iconic counts we found that our above-described elegant wave count in gold had been replaced by a count that's much less attractive and has a lower score. We would then dig into EWAVES' abundant output to learn exactly what transpired and use that information to keep, reject or amend our adjustment. More about that in an upcoming issue.

**What's next?**

At this point, the foundation of EWAVES 2 is complete, along with an extensive array of internal tests to maintain the functionality of the software. Now we're working on the meat of the program and on improving its trading strategies. We'll keep you posted.

## FREE TO SUBSCRIBE, FREE TO SHARE

The *EWAVES Flash* publication is open-access. Feel free to share the link ewaves.com/1404EWF for this issue, or visit ewaves.com and click on the "newsletter" tab to read other issues or to sign up for alerts when new issues are published. To subscribe to Flash Services, use the "services" tab or go directly to ewaves.com/flash-services

## CITATIONS

[1] Forward engineering object recognition : a scalable approach. (n.d.). DSpace@MIT:. Retrieved April 10, 2014, from http://dspace.mit.edu/handle/1721.1/62622

[2] Carruthers, S., & Stege, U. (2013). On Evaluating Human Problem Solving of Computationally Hard Problems. The Journal of Problem Solving, 5(2), 45.

[3] Memory part 2: CPU caches. (n.d.). [LWN.net]. Retrieved April 10, 2014, from http://lwn.net/Articles/252125/

*EWAVES Flash* is published by Elliott Wave International, Inc. Mailing address: P.O. Box 1618, Gainesville, Georgia, 30503, U.S.A. Phone: 770-536-0309. All contents copyright © 2014 Elliott Wave International, Inc. The *EWAVES Flash* publication may be freely shared and redistributed for non-commercial purposes provided it is not modified from its original form and all credit is given to Elliott Wave International, Inc. You may not use this work for commercial purposes without the permission of Elliott Wave International, Inc. Otherwise, feel free to quote, cite or review if full credit is given. Dissemination of Flash recommendations made by Flash Services is strictly forbidden.

FINANCIAL NEWSLETTERS INTERNATIONAL FOUNDING MEMBER

**Contact EWI Customer Care:** Call 770-536-0309 (internationally) or 800-336-1618 (within the U.S.). Or email customercare@elliottwave.com.