My Family and Other Algorithms

Dave Cliff

(A short invited article for the 2007 Euromoney Algorithmic Trading Handbook.)

In the summer of 1996, I invented a trading algorithm that outperforms human traders. The algorithm, known by the acronym "ZIP" (for Zero-Intelligence-Plus) is embarrassingly simple. But being very simple makes ZIP very fast, which, nowadays, with major players running their algorithmic trading operations on proximity-servers huddled as close as physically possible to the main servers on major exchanges, is another point in ZIP's favour. One other thing that makes ZIP attractive to many people is that you can use it for free (and I've been told that some major algo houses do exactly that): full details of the algorithm, including the "C" source-code, were published on the web in 1997 and have been available there ever since.

In the ten years since it was invented, other algorithms have come along to challenge ZIP; but over the same period I've been involved in the development of various extended versions of the original ZIP: there is now a whole family of ZIP algorithms, and the original has been re-christened "ZIP8" (for reasons that will become clear later in this article). This family of extended ZIPs give significant improvements in performance, with little or no increase in latency. The most recent of the ZIPs to be given away for free goes by the name of "ZIP60". I'll come back to ZIP60 later below.

Before that, I'll explain how ZIP came to be invented, and the basics of how it works.

In 1996 I was working as a lecturer in computer science and artificial intelligence at the University of Sussex, and supplementing my meagre academic salary by day-trading equity options on LIFFE, the main London futures exchange. The good folk at Hewlett-Packard Labs in Bristol offered me a seven-month stint working with them as a visiting academic, and told me that I could come up with my own research project for my stay there. In those days, trading at LIFFE was still based heavily on face-to-face open-outcry trading pits populated by human traders, but I thought it would be scientifically interesting to come up with some software that could, in principle, do the job of a human trader in the pit. The transition to screen-based trading that was by then complete in the London equity market had to happen to LIFFE eventually, I reasoned, so maybe one day I could use my algorithms to do my options trading for me. On that basis, I proposed to HP that I'd develop algorithms for trading in electronic auctions (in the language of academic economics, an *auction* is any mechanism – expressed as a set of rules and processes – by which buyers and sellers interact to agree a price for a transaction). Almost all traders in the global international financial markets interact via a particular form of auction market mechanism known as the *continuous double auction* (CDA), where buyers quote increasing bid-prices and sellers quote decreasing offer-prices, with any trader free to accept any counterparty's quote at any time, and without any synchronization or central auctioneer. The CDA is the mechanism that operates in open-outcry trading pits, and is also the basis of most screen-based trading systems, in all the world's major financial markets: equities, FX, commodities, and derivatives. HP were very happy with this proposal. My employer, Sussex University, imposed one constraint: they wanted my work published as an HP Labs technical report, because such publications are the output by which academics are judged. It was this that led to ZIP being given away for free; and ZIP being freely available played a significant role in it becoming widely known since.

The basic assumption behind ZIP is that at any one time a trader is either buying or selling, and that either way the trader has a private (i.e., secret) limit price, denoted by the Greek character λ, for a unit of the good being traded. For this discussion, we'll consider the simplest case: where each trader is just looking either to buy or sell one unit, and so the trader has only one λ value. (ZIP does work with traders buying and selling multiple units, and it has also been used for traders that arbitrage by buying from one market and selling into another). For a buyer, λ is the maximum price that the trader is prepared to pay; for a seller, λ is the minimum price that the trader is prepared to sell at. Of course, wandering into a CDA and making your first quote be λ is the dumbest strategy of all: it leaves you no room to negotiate. So, ZIP also assumes that each trader keeps λ secret and instead quotes a price p in the market that is greater than λ (for sellers) or less than λ (for buyers), where the difference between λ and p is determined by that trader's *margin*, denoted by the Greek character m.

At the core of the ZIP family of algorithms is a minimally-simplistic set of qualitative heuristics (that is, some rough rules-of-thumb) for adjusting trader i's margin m at time t – denoted by $m_i(t)$. These heuristics depend on four things: (1) i's current quote-price $p_i(t)$; (2) whether i is still active in the market (i.e., still attempting to trade, or just watching); (3) whether the last price quoted in the marketplace $q(t-1)$ was an *offer* or a *bid*; and (4) whether $q(t-1)$ was *accepted* by a counterparty or not. The ZIP heuristics were arrived at by introspection (that is, I asked myself what *I* would do if it was *me* doing the trading) followed by some trial-and-error refinement. Nevertheless these heuristics have been found to work extraordinarily well in practice. The heuristics for any one trader can be thought of as defining a simple "decision tree" – a sequence of yes/no questions – leading to outcomes that specify whether the trader should raise or lower its margin, or just leave the margin unchanged. For the cases where a change in the margin is required, a process for adjusting or *adapting* the trader's margin parameter is specified, as a set of equations that easily fit on one page: the equations wouldn't cause any trouble to a first-year university student of mathematics. Each ZIP trader has its own pair of parameters that govern the way in which it adapts $m_i(t)$ in reaction to market events. These parameters are the trader's learning rate b_i (which determines how fast it adjusts m_i) and its "momentum" g_i (which determines how volatility is smoothed over). And that is pretty much all there is to it. I designed this algorithm while I was at HP, wrote the required technical report, and then headed back to academia. I carried on tinkering with ZIP, and when I moved to MIT in 1997 I had a few students work on exploring some ideas I'd had, plus I developed a way of having the computer automatically optimize the values of the various parameters in a ZIP trader, adapting it to particular market conditions. In the original ZIP, there were 8 such key parameters, which is why that ZIP is now known as ZIP8. The automatically-optimized ZIP8 traders were better than the hand-designed ones that I had originally created at HP, so ever since then the newer additions to the ZIP family of algorithms has assumed that a computerized process will search for useful values of the various control parameters. Once you've accepted that a computer will be choosing the parameter-values for you, there's no need to keep the number of control-parameters small enough to make choosing their values simple enough for humans. So, ZIP60 gets its name from the fact that it has 60 control parameters. The extra parameters give it greater sophistication than ZIP8, and its adaptation and responses can be significantly more fine-tuned to the prevailing market conditions.

Tiring of life as an academic, I quit MIT to join HP Labs as a research scientist. A few years later, in some surprising experiments done at the research labs of HP's rival IBM, it was shown that ZIP traders (and also IBM's own "MGD" trading algorithm) consistently out-perform human traders in human-against-robot screen-based CDA markets. In the IBM experiments, the robot traders consistently made profits a few percentage points higher than did the human traders they were competing against, and ZIP had the joint-highest average efficiency of the algorithms used in the IBM study (roughly, "efficiency" measures how much of the maximum potential margin is actually achieved by the traders: higher efficiency is more desirable). Mean values calculated from the efficiency data presented in the IBM results are 1.030 for ZIP (mean from n=2 experiments); 1.023 for MGD (n=4); and 0.876 for humans (n=6).

In discussing the possible impact of their work, the IBM team wrote that the "…successful demonstration of machine superiority in the CDA … could have a … powerful financial impact – one that might be measured in billions of dollars annually", and in their conclusions they speculate on the future possibility of online e-marketplaces currently populated by human traders becoming populated entirely by trader agents.

Those results got a lot of press coverage world-wide, and lots more people downloaded the ZIP technical report and started using it. I ended up working as a consultant on HP's research projects with major exchange operators and investment banks; then I moved to Deutsche Bank; and then back into academia -- from where I've been able to get involved in some entrepreneurial ventures centered on automated trading systems. To the best of my knowledge, ZIP and MGD remain the only two algorithms shown to have consistently outperformed human traders in rigorous experimental conditions.

ZIP is an algorithm for automated execution of trades: it doesn't do prediction or make investment decisions. But, having been involved in frontline research on Algorithmic Trading for over a decade now, I think it is pretty obviously true that the day on which human traders are acknowledged as being too slow, to error-prone, and too limited in bandwidth and in processing power, is drawing ever closer ever faster. An inconvenient truth for many traders, but a truth nonetheless.

This article is necessarily short. For further details, see the ZIP-trader home-page at www.ziptrader.org. On that web-page there are links to the old ZIP work at HP, to the papers by the IBM team, and to the latest papers on ZIP60. That's also the place where the latest and most sophisticated evolution of ZIP, better even than ZIP60, will be made public sometime in the coming months. Watch that space.

Contact Info, until June 30th 2007:
School of Electronics and Computer Science, University of Southampton, SO17 1BJ, U.K.
Telephone +44 23 8059 4510 or +44 7977 552250
Fax: +44 23 8059 9179
Web: www.sense.ecs.soton.ac.uk or www.syritta.com
Email: dc@ecs.soton.ac.uk or davec@syritta.com

Contact Info, after July 1st 2007:
Department of Computer Science, University of Bristol, BS8 1UB, U.K.
Telephone +44 117 9545264 or +44 7977 552250
Fax: +44 117 9545208
Web: www.cs.bris.ac.uk or www.syritta.com
Email: dc@cs.bris.ac.uk or davec@syritta.com

[Note: as publication occurs after 1st July 2007, Bristol-based bio is given below.]

Biography. Dave Cliff is Professor of Computer Science at the University of Bristol, UK, and owner of Syritta Algorithmics Ltd, a trading technology research and consultancy company. Previously he has held faculty positions in the UK at Sussex and Southampton Universities, and in the US at the MIT Artificial Intelligence Lab. He has also worked as a research scientist for Hewlett-Packard Labs Europe, and for Deutsche Bank London's Foreign Exchange Complex Risk Group.