

March 16th, 2016

Christopher J. Kirkpatrick
Secretary of the Commission
Commodity Futures Trading Commission
Three Lafayette Centre
1155 21st Street, NW
Washington, D.C., 20581

Re: Proposed Rule, Regulation Automated Trading (“Regulation AT”) RIN 3038-AD52

Dear Mr. Kirkpatrick,

We thank the Commission for this opportunity to comment on the proposed Reg AT (the Proposal).

In this comment letter, we would like to bring your attention to recent advances in the field of formal verification that will have tremendous impact on the safety and fairness of financial markets. Formal verification, the scientific discipline dedicated to the mathematical analysis of algorithms, is already relied upon by safety critical industries such as avionics and microprocessor design. We believe that many of the goals of Reg AT can be met by appropriate application of formal verification.

In the context of Reg AT, we would like to highlight how formal verification can provide:

- rigorous, quantitative approaches to system testing (“model-based testing”),
- mathematically-precise communication of client (and counter-party)-facing algorithm specifications,
- automated analysis of algorithm specifications for many key regulatory properties.

Our views in this letter are representative of our broader ‘push’ for the financial industry to adopt formal verification for the design, implementation and regulation of algorithms. Our company website (www.aestheticintegration.com) contains further publications and information on our mission.

About Aesthetic Integration

Aesthetic Integration Ltd. (AI) is a financial technology start-up bringing cutting edge formal verification technology to financial markets. AI is working with leading financial institutions to revolutionize the process of designing, implementing and ensuring compliance of complex trading systems. In late 2015, AI won the UBS Future of Finance Challenge, coming in 1st out of 620 companies from 52 countries. AI also won a Futures Industry Association Innovator Award at the FIA Expo 2015 in Chicago.

Denis Ignatovich, co-founder of AI, has over a decade of experience in trading, risk management, quantitative modeling and complex trading system design. Prior to joining AI, he was head of the central risk trading desk at Deutsche Bank London. He holds an MSc in Finance from the London School of Economics and undergraduate degrees in Computer Sciences and Finance from the University of Texas at Austin.

Dr. Grant Passmore, co-founder of AI, has more than ten years' industrial formal verification experience. He has been a key contributor to safety verification of algorithms at Cambridge, Carnegie Mellon, Edinburgh, Microsoft Research and SRI. He earned his PhD in Automated Theorem Proving from the University of Edinburgh and is a Life Member of Clare Hall, University of Cambridge.

Testing

Many (perhaps all) recent financial algorithm regulations profess the need for “thorough testing.” Furthermore, the Proposal makes numerous references to recent advisory notes and directives and their testing requirements. However, nowhere does the Proposal (nor the referenced documents) define “thorough testing.” While we appreciate that the Commission is taking a principles-based approach to Reg AT, some fundamental questions must still be addressed:

- What is thorough testing?
- How does one ensure that testing done is appropriate for the complexity of the system being tested?
- What metrics (i.e., quantitative measures) should be used to assess the thoroughness of testing?

In Question 43, the Proposal asks “Are the procedures described above for the development and testing of Algorithmic Trading sufficient to ensure that algorithmic systems are thoroughly tested before being used in production, and will operate in the manner intended in the production environment?” We believe the answer is no, as the fundamental questions above remain unanswered. Without addressing these questions, how can the sufficiency of a firm’s testing methodologies be assessed?

For example, consider the ‘glitch’ (“Algorithmic Event”) that caused BATS to fail during its own IPO several years ago. BATS has a considerable share of the overall market and, in many ways, represents industry best practices. In the words of the exchange’s CEO, the glitch resulted from “[a] combination [that] hadn't been seen in the hundreds of tests we'd run before this”¹. As an industry, we must recognize that the currently employed approaches to system testing are insufficient.

A number of comments referenced in the Proposal highlight the availability of DCM historical data for backtesting. Although these tools may be useful for understanding market risks and P&L of trading strategies, they are, by definition, insufficient for the elimination of unforeseen Algorithmic Events. There are countless recent examples where the assumption that “the future resembles the past” has led to disaster.

Consider the microprocessor industry. Before Intel released their flawed Pentium microprocessor in 1994, they had performed millions of tests on its floating point unit. The incorrect behavior was not caught by any of these tests. After a recall (costing Intel nearly \$500M), their answer was not to simply run more tests. Instead, they adopted the rigorous use of formal verification to mathematically analyze their floating point

¹ <http://www.wsj.com/articles/SB10001424052702303404704577304034248567486>

designs before they are released. Similar techniques can be brought to bear on financial algorithms, both for analyzing specifications and for the generation of test suites with rigorous, quantitative coverage metrics.

The Precise Specification Standard

We are encouraged by your statement that the venue² operators should “clearly communicate such policies and procedures to market participants” when referring to venue operational details. We argue the Commission should go further in requiring the operational details made publicly available be analyzable by modern automated reasoning techniques. Such formatting will allow the Commission and the market participants to tap into the field of formal verification, in a manner similar to how the FAA and DoD analyze and regulate complex, safety-critical algorithms. More specifically, such format will allow those trading on the DCMs to automatically analyze their algorithms’ consistency with the specifics of venue algorithms and compliance with regulatory directives.

In particular, the Proposal’s Question 76, asks *“The Commission proposes that DCMs provide a description of the relevant material attributes in a single document “disclosed prominently and clearly” on the exchange’s website. The Commission also proposes that this document be written in “plain English” to allow market participants, even those not technically proficient, to understand the attributes described. Would these requirements be practical and help market participants locate and understand the information provided?”*

We argue that ‘plain English’ is the wrong format for communicating operational details of a DCM (e.g. order type definitions, connectivity protocols and matching rules). We see the following issues when a DCM operator only makes available an English Description (ED) of the operational details of the venue:

1. An ED is completely detached from the actual production system. There is no way to automatically check an ED’s correctness, to understand if it properly describes the trading system. So, if a member of the development team makes a subtle, yet significant change to the logic of the system (perhaps accidental), then the ‘break’ with the ED is very difficult to detect; checks for this are not automated.
2. An ED is typically describing a system that may be in an infinite (or virtually infinite) number of possible configurations. In practice, the possible behaviors of such a system cannot be exhaustively analyzed “by hand,” i.e., by a team of regulators and market participants reading and attempting to understand the ED.³
3. An ED cannot be used to test implementations of systems that are trading on the venue described. In other words, the ED cannot be incorporated into the development process of trading systems that send orders to the venue.

In order to remedy the issues above, we suggest the following core requirements on any regulatory-compliant format for disclosing operational details of a DCM. We call a disclosure format meeting these requirements a Precise Specification (PS).

1. In a PS, the operational details of a DCM should be described in an executable programming language with a formal semantics. We give detailed examples of this in our white papers. The phrase ‘formal

² We use the term ‘venue’ interchangeably with the terms ‘exchange’ and ‘DCM’.

³ Please see our white paper “Case Study: 2015 SEC Fine Against UBS ATS” for a concrete example.

semantics' refers to the ability to translate the algorithm disclosed into a precise mathematical model. This model can be analyzed using formal verification techniques to automatically check its logic for potential violations of many key regulatory directives.

2. Questions may arise as to whether certain information pertaining to the actual operation of a DCM should be disclosed. We propose a simple test that should help the venue operators and the Commission answer that question:

Disclosure Test: *Is the information necessary for one to write an observationally-equivalent simulator of the venue? If so, it should be required.*

In a PS, the Disclosure Test must be passed. One byproduct of this is that, given a PS, one can directly create an executable "simulator" of the venue design that was disclosed. Different levels of observational-equivalence can be utilized for different contexts and purposes. As a base-line, we suggest the observational-equivalence of two venue implementations require equality of message sequences (FIX,ITCH, proprietary binary, etc.) over 'replays.'

An operator should judge the disclosed PS complete if and only if it is reconcilable with actual post-trade data of the production system. That is, the specification given should be convertible into a machine-executable program that can be run against actual historical data. In doing so, its faithfulness to the behavior of the production system may either be confirmed or refuted over a given time window.

Given a PS, the Commission and financial firms can leverage formal verification techniques to:

1. Automatically analyze the submitted specifications for potential violations of regulatory directives and design inconsistencies. For example:
 - Is the order ranking criteria transitive?
 - Is there any subtle combination of order parameters that allow someone to 'jump the queue'?
 - Is the venue designed to handle different market regimes?
2. Allow those trading on the venues to automatically test their connectivity and verify their routing algorithms using model-based, quantitative state-space coverage metrics (discussed below).

For more specific examples of PS formats, we urge the Commission to review our recent white papers detailing an application of our formal verification system, Imandra, to analyzing safety and fairness properties of venues. In fact, our latest white paper "Case Study: 2015 SEC fine against UBS ATS" follows in detail the SEC's order against UBS ATS. The case study formed part of our application into UBS's Future of Finance Challenge that was held in (August - December) 2015. Aesthetic Integration was selected as the first place winner out of 620 applicants from 52 countries. In the case study, we detail how issues raised in the SEC settlement, namely sub-penny pricing and undisclosed trading constraints, may be encoded as mathematical properties and automatically analyzed for a venue specification.

About Formal Verification

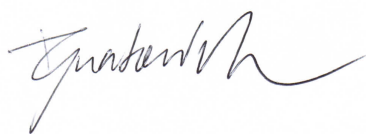
We published a white paper in 2015, ‘Creating Safe and Fair Markets’⁴, describing formal verification, how it is currently applied to other industries, and the recent advances that power our application of formal verification to financial markets. In summary, formal verification is an interdisciplinary field of mathematics, computer science and artificial intelligence directed towards analyzing the behavior and implementation of complex algorithms. It is widely relied upon within the US federal government. To list a few examples:

- The FAA requires⁵ precise levels of system testing and formal verification within both the Common Criteria Evaluation Assurance Levels and DO-178C⁶ frameworks. Safety-critical algorithms such as air traffic control, onboard autopilots and collision avoidance, and the security of aircraft local area networks must satisfy these rigorous requirements before they are allowed to be deployed.
- The Department of Transportation has commissioned work⁷ on creating a formal verification framework for regulating the safety of autopilot algorithms inside self-driving cars and other autonomous vehicles.
- NASA is one of the biggest drivers in the field. Among many other high-profile examples (Mars rovers, etc.), NASA’s NextGen Air Traffic Management⁸ framework relies on formal verification to ensure its safety.
- The Department of Defense⁹ leverages formal verification across numerous applications, including the design and regulation of cryptographic algorithms and secure hypervisors.

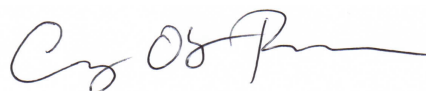
Concluding Remarks

Again, we thank the Commission for this opportunity to express our comments on the proposed Reg AT. We hope you find our comments useful.

Sincerely,



Denis Ignatovich
Co-Founder, AI



Grant Passmore, PhD
Co-Founder, AI

⁴ “Creating Safe and Fair Markets” is available from www.aestheticintegration.com

⁵ See <https://buildsecurityin.us-cert.gov/articles/best-practices/requirements-engineering/the-common-criteria>

⁶ See http://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_20-115C.pdf

⁷ See <http://utc.ices.cmu.edu/utc/utc-tset-projects.html>

⁸ See <http://www.hq.nasa.gov/office/aero/asp/airspace/>

⁹ See <http://www.darpa.mil/program/high-assurance-cyber-military-systems>