


Distributed system for algorithmic trading

Michal Hornický



Abstract

The success of cryptocurrencies like Bitcoin has created many new opportunities. One of them came somewhere around the year 2012-2013, in a form of an online cryptocurrency exchange. Since then, many new online exchanges were created. These exchanges provide unprecedented ease of use and access to everyone, contrasting existing financial exchanges. Day-trading*on these exchanges is easy, and has a large potential because of the extreme volatility of these new markets. This paper outlines the design and implementation of a distributed system, that would facilitate this task. The goals, which include ease of use for new users, scalability for large number of users, and customization for advanced users, combined with problem domain pose interesting requirements, which influenced the design and implementation.

Keywords: Automated trading — Distributed systems — Rust

Supplementary Material: [Github Repository](#)

*xhorni14@fit.vutbr.cz, Faculty of Information Technology, Brno University of Technology

1. Introduction

[Motivation] The innovations in financial sector, mainly cryptocurrencies like Bitcoin, have created new opportunities. One of them, is the arrival of multiple online exchanges, that focus on cryptocurrency trading. These exchanges have very low barrier to entry, and can be used for easy day trading. Goal of this project is creation of a website, that would allow automatization of day trading on these exchanges.

[Problem definition] In order to effectively support multiple users, the designed system must be able to seamlessly scale according to computing load. It must allow its users to easily create multiple trading strategies, and execute trades on cryptocurrency exchanges.

[Existing solutions] Most attempts to fail at one or more of the requirements. The older solutions are mostly command-line applications that require complicated installations, and hardware that must be continu-

ously managed. One of these is the **Gekko**¹ trading bot. Main drawback of this solution is the use of JavaScript as the implementation language, and the requirement of Node.js.

Other cloud based solutions, that are more closely related to our approach are usually overly complex, requiring user to write complex strategies that must decide not only when to execute trades, but specifics of these trades. One of these systems is **CryptoTrader**², that uses CoffeeScript as a language for implementing user strategies. This system supports multiple assets within one strategy, making them extremely expressive. However, the drawback of this approach is the increased complexity.

[Our solution] our solution aims to surpass other automated trading systems in several aspects. Thanks to the decision to implement the system as a web ap-

¹<https://gekko.wizb.it/>

²<https://cryptotrader.org/>

37 plication, we remove all local software requirements,
38 making the system very approachable. Thanks to the
39 distributed architecture, the system will have the nec-
40 essary degree of scalability.

41 [Contributions] Implemented system is built upon
42 scalable architecture that utilizes cloud environment,
43 is able to scale from single to thousands of users seam-
44 lessly, It also allows sub-second latency between re-
45 ceiving of new financial information and possible exe-
46 cution of actual trades on real exchange. Implemented
47 system currently utilizes only one exchange, however
48 support of additional exchanges should be extremely
49 easy.

50 2. Theoretical background

51 In order to understand automated trading systems, we
52 must first understand how, modern exchanges operate.
53 The core concept of an exchange is the price discovery
54 mechanism. In short, this means, that the exchange
55 does not determine the price of an asset, but rather
56 the price is "discovered" by interactions of individual
57 actors on the exchange. In simplistic terms, this corre-
58 sponds to supply-demand market mechanism. When
59 the supply of an asset is larger than demand, the price
60 falls, and when the demand rises, the price rises ac-
61 cordingly.

62 2.1 Automated trading

63 Today, most of the trading performed even on con-
64 ventional exchanges is done by automated systems.
65 Origins of these systems can be traced to the 1980s,
66 but probably the biggest milestone was when IBM
67 in 2001[1] experimented with automated trading, and
68 implemented system consistently outperformed even
69 professional traders.

70 2.2 High frequency trading

71 Modern incarnation of high-end automated trading sys-
72 tems is called High Frequency Trading(HFT). These
73 systems are commonly co-located with the exchanges,
74 aiming for lowest possible latency between receiving
75 financial data, and execution of market orders. We can
76 divide them into several groups based on the decision
77 process used for creating market orders.

78 We will focus on **Tick-data market making** strate-
79 gies. These strategies that utilize periodic information
80 the about price of an asset in order to determine short
81 and long term trends of this price. Based on the short
82 and long term trends these strategies forecast the price
83 into the future.

84 The benefit of this approach is mainly simplicity.
85 These strategies do not have to rely on complex data

describing real world events, that might influence price
of an asset(eg. company mergers), and they do not
have to perform actual trades explicitly.

2.3 Computing environment

While strategies outlined earlier are easy to implement,
they require non-trivial amount of computing power.
Coupled with the the need to support multiple users,
the requirements for computing power needed to run
this system grow.

In order to provide this amount of computing power,
we decided to design and implement the system using
distributed architecture. This means, that the system is
written in way, that allows its individual components
to operate separately, and be deployed on different
machines. To achieve this goal, we have chosen to
use Cloud computing approach³ on the deployment
side. And utilize Actor Model as the core paradigm on
implementation side.

2.4 Actor model

Actor model is a conceptual model of describing con-
current computation[2]. Each actor can: Create new
actors, send messages, modify its state and decide how
to respond to received messages. Primary constraint of
this model is the restriction of modifying application
state. Each actor can modify its local state however
it wants, but can only affect other actors by sending
messages.

2.5 Rust & Actix

Due to the choice of the Actor model as a core paradigm,
the choice of possible implementation languages was
limited. The chosen language would have support this
programming model (either implicitly, or through the
use of a library). Other considerations included were
the runtime overhead, safety, ease of integration with
other technologies. Evaluated languages and frame-
works include: C# with Akka.NET, Erlang, Java with
Akka, and Rust with Actix.

Rust with the Actix library was chosen mainly
due to extremely low overhead of this programming
language (not requiring a VM), ease of integration
with other technologies (LUA), and due the authors
personal interest in this language, and corresponding
library.

3. Designed system

Probably the biggest obstacle to the implementation
of the system was its distributed nature. On the imple-
mentation side, this meant the use of Actor model, as

³https://en.wikipedia.org/wiki/Cloud_computing

133 a core architectural paradigm. The use of the Actix
134 library has simplified many challenges with the use of
135 this computing paradigm, but it also came with some
136 drawbacks. The library allows seamless use of actors
137 within single or multi-threaded environments, support-
138 ing use of single or multiple concurrent threads, but it
139 does not contain an implementation of primitives that
140 would allow actors to communicate between processes
141 or even different machines.

142 3.1 Communication

143 Therefore, part of this project was the design, and
144 implementation of this capability. We have designed
145 and implemented the **actix-zmq** library, that provides
146 actors for communication over the ZeroMQ network-
147 ing technology, and **actix-comm** library that provides
148 abstractions for implementing simple Request-Reply
149 services, Publish-subscribe pipelines, and other sup-
150plementary components (eg. Load balancing broker
151 for services). The **actix-comm** library builds on top of
152 **actix-zmq**, and both of them should be usable in other
153 projects, and will be published as separate libraries,
154 that should enrich already rich ecosystem around the
155 **Actix** library.

156 3.2 Deployment

157 Since we are using Cloud environment as our primary
158 deployment target, this side of the system also had
159 to be adapted. We decided to use **Kubernetes** as a
160 primary tool for managing our deployments.

161 Kubernetes is an orchestration tool, used for au-
162tomated deployment and management of distributed
163 systems running in the cloud environment. Kubernetes
164 defines a set of primitives, which are used to describe
165 a distributed system. The kubernetes runtime then
166 dynamically modifies state of the system, to conform
167 to described model. The kubernetes runtime runs on
168 a Cluster. A cluster is comprised of multiple virtual
169 machines(Nodes), and can dynamically scale number
170 of used nodes.

171 3.3 Actual system

172 The actual system is then designed as a set of loosely
173 coupled components. Each component is comprised of
174 several kubernetes Pods, managed by a Deployment,
175 and exposed by a Service. Within each pod, there
176 might be multiple containers, but most of them only
177 use single one.

178 Here are components that that describe our system
179 in simplest terms

- 180 • Exchange - provides interface to a specific ex-
181 change, currently only the Bitfinex exchange,

- Core - Receives updates from exchanges, De- 182
cides when to evaluate strategies against this 183
data, and forwards trading decisions to individ- 184
ual exchanges. 185
- Eval - Evaluates strategies using multiple load- 186
balanced workers 187
- Web - Provides web interface for user interaction 188
- Storage - Stores financial and user data. 189

4. Implementation 190

As mentioned earlier, the implementation was per- 191
formed using the Rust language on top of Actix actors 192
as a basic architectural blocks. It is currently divided 193
into 2 executables, The `web` executable houses the user 194
interface implemented using `actix-web` as a back-end, 195
and `LitElement`⁴ based web application as front-end. 196

The second executable is the `trader` application. 197
This is implemented as a command line application, 198
that contains the implementations of several different 199
components, and should be split into separate exe- 200
cutable for each component in the future. 201

4.1 Data flow 202

The whole system is best described by the type of data 203
it consumes, and how this data flows throughout it. 204
Primary data sources are individual exchanges, and the 205
web application. The web application only communi- 206
cates with the database, and thus is not that interesting 207
in this aspect. 208

However, the exchanges are more interesting. Most 209
cryptocurrency exchanges provide REST API used for 210
executing trades, and WebSocket endpoint, that is used 211
for providing latest financial data. For each exchange 212
supported by the system there is a dedicated compo- 213
nent, that serves as an adapter to this exchange. Main 214
purpose of an exchange adapter is translation system 215
requests into a specific exchange API requests, and 216
forwarding the updates received over WebSocket to 217
core system component. 218

The individual exchange adapters each connect 219
using **PUB** ZeroMQ socket to core service. This forms 220
a Fan-In topology, that would be difficult to implement 221
using other technologies. 222

The data flowing from exchanges is in the form of 223
per-minute OHLC⁵ data. This is then processed by 224
the core component, which removes duplicate entries, 225
computes data points for different time-scales, and 226
publishes them along with received updates. During 227
this step, the data is also stored into persistent stor- 228

⁴<https://lit-element.polymer-project.org/>

⁵https://en.wikipedia.org/wiki/Open-high-low-close_chart

229 age, which currently takes the form of a PostgreSQL
230 database.

231 The decision actor in core component periodically
232 loads the information about assignment of strategies
233 to individual assets. Whenever it receives new OHLC
234 data it determines which strategies should be evaluated,
235 and sends this information to the evaluation compo-
236 nent, which is implemented as a load-balancing broker,
237 with multiple workers.

238 Whenever an evaluation worker receives an evalua-
239 tion request, it retrieves the strategy text, and historical
240 data from the database. It then creates a new Lua VM,
241 configures it (eg. disabling file access), and provides a
242 suite of analytical functions, that can be used by indi-
243 vidual strategies. Then it loads the strategy script into
244 this VM, and executes it.

```
// Simple moving average
local sma = ta.sma(10)
// Exponential moving average
local em = ta.ema(10)

// Short term > long term
if sma() > ema() then
    return "short"
else
    return "long"
end
```

Figure 1. Example strategy

245 The output of the strategy is the desired market
246 position - "long" or "short", the former denoting own-
247 ership of target asset and the latter denoting the owner-
248 ship of the exchange currency, eg. US Dollars.

249 This work does not focus upon the individual strate-
250 gies, or methodologies behind them, it only provides
251 basic building blocks for creating them. One of future
252 enhancements might the support of more advanced
253 types of strategies.

254 Then upon receiving the results of strategy eval-
255 uation, the core component checks whether there is
256 trading account information associated with the as-
257 set. If there is, it then sends a request to an exchange
258 adapter, which then might check the current market
259 position of the user, and possibly execute one or more
260 trades, ensuring that the requested market position is
261 achieved on this trading account.

262 5. Conclusions

263 **[Paper Summary]** This paper outlined the conceptual
264 idea behind the project, the issues encountered and

how they influenced the design and an actual imple- 265
mentation of the system. The system is implemented 266
as a distributed application, with focus on scalability, 267
and is accessible using a web application, satisfying 268
the usability requirements. 269

[Highlights of Results] The implemented system 270
currently supports single exchange, and over 200 dif- 271
ferent asset pairs. Each of these asset pairs can have 272
a single LUA strategy, and single trading account as- 273
sociated with it. The system supports executing large 274
number of strategies, with sub-second latency between 275
updates from an exchange, and execution of trades on 276
these exchange. Compared to command-line applica- 277
tion solutions, our system can support arbitrary num- 278
ber of strategies, with possible future improvement of 279
tracking individual strategy performance. Compared 280
to other cloud based solutions, our system provides 281
extremely easy strategy implementations 282

[Paper Contributions] Achievement of these goals 283
was mainly possible due to distributed approach. How- 284
ever, this approach brought its own set of complica- 285
tions, which required solutions. These solutions were 286
implemented in support libraries `actix-zmq` and 287
`actix-comm`, that should be useful in other projects 288
with similar goals. 289

[Future Work] While functional, the system lacks 290
several pieces of advanced functionality (eg. More 291
complex strategies or testing strategies on historical 292
data). The system should be also extended, to support 293
multiple exchanges. In addition to extending the actual 294
system, the support libraries mentioned earlier are also 295
good targets for future development. 296

Acknowledgements 297

I would like to thank my supervisor RNDr. Marek 298
Rychlý Ph.D. for his help, and valuable advice regard- 299
ing this project provided during frequent meetings. 300

References 301

- [1] Gerald Tesauro and Rajarshi Das. High- 302
performance bidding agents for the continuous 303
double auction. In *Proceedings of the 3rd ACM* 304
Conference on Electronic Commerce, EC '01, 305
pages 206–209, New York, NY, USA, 2001. ACM. 306
- [2] Carl Hewitt. Actor model for discretionary, adap- 307
tive concurrency. *CoRR*, abs/1008.1459, 2010. 308