LONG FINANCE

DISTRIBUTED FUTURES

Timestamping Smart Ledgers

Comparable, Universal, Traceable, Immune

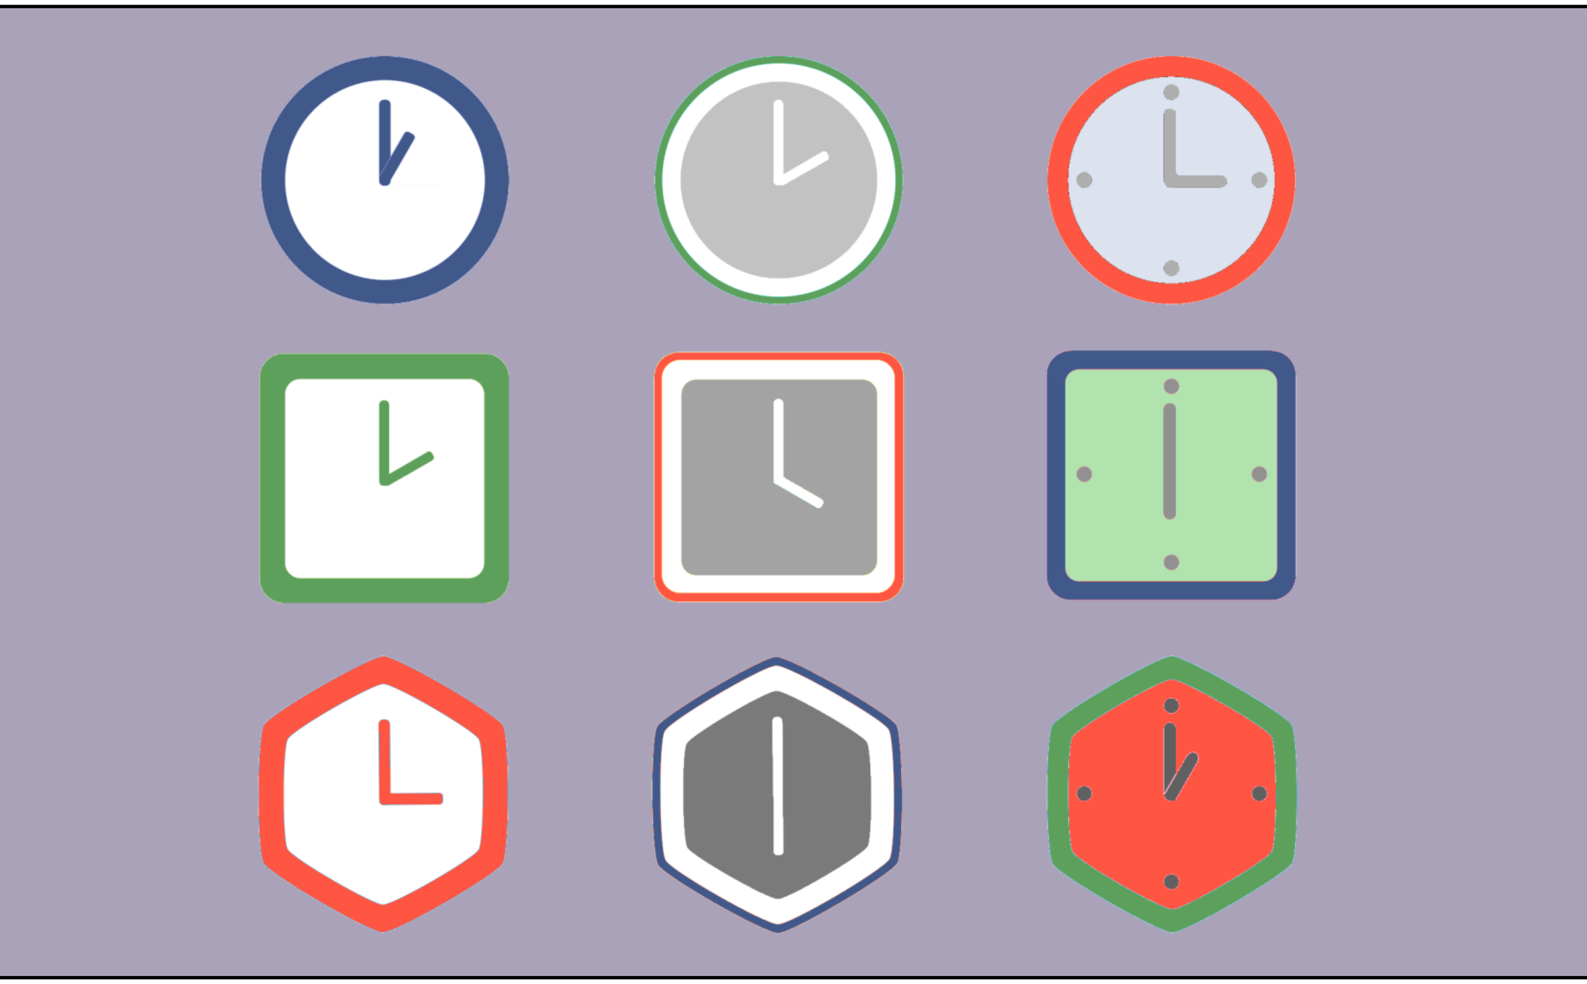

June 2018

CARDANO FOUNDATION

The Z/Yen Group

Timestamping Smart Ledgers

Comparable, Universal, Traceable, Immune

June 2018

Sam Carter
Financial Sector Researcher and Quant Developer

Foreword

What time is it? When did something happen? When is an event expected? How much time separates two events? Which came first?

These questions arise in everyday life and in every kind of discipline. For centuries clockmakers have striven to improve ways to measure time and to achieve ever higher degrees of precision, to meet the demand for its closer and closer estimation – and by achieving that improvement have enabled ever more exacting applications of timekeeping, as the measurement of time has become almost unimaginably precise. Over that period changes in technology have also allowed a uniform time to be widely known to an accuracy not just to the hour but to the minute and then to the second and within even smaller tolerances, in public, in our houses and places of work, and carried on our persons. Centralised time synchronises our local time, and sub-microsecond accuracy routinely enables satellite navigation – relying on the connection of 'when' with 'where' by the speed of light.

In recent years, computing has transformed the speed of financial trading, and following this, the recording of trades – and in particular when they took place – has become far more stringently regulated, with timestamping to the microsecond required in some areas. The story of John Harrison and his marine timekeepers is an example of how an eighteenth-century authority, the Board of Longitude, required demonstration that his new technology was sufficiently accurate and fit for purpose, navigation at sea. Likewise, the idea of using the novel technology of Smart Ledgers for keeping accurately time-stamped records, particularly in the financial world, will need to be examined by present-day regulators and demonstrated to be suitable to meet their requirements. The sequential nature of Smart Ledgers is a valuable inherent property, but interesting challenges of incontrovertibly recording time remain. So, I welcome this research that explores accuracy AND precision in the application of time in the brave new world of Smart Ledgers and the impact they will have on all of us, over time …

Andrew C. James.

Andrew James
Master, The Worshipful Company of Clockmakers

Table Of Contents

The front cover is designed by Alexandra Karathodorou and is begging the answer to the question: Which clock is the odd one out? (See page 9 for the answer).

Preface

Clocks and their bells and towers were enormous social projects, proliferating throughout Europe from about 1270. Salisbury Cathedral has hosted a clock since 1306 and its current one, dating from 1386, is the oldest working clock in England. Cockneys are those born within the sound of the Bow bells and London has the “Oranges and Lemons” rhyme for several church bells. London’s oldest business until 2017 was the Whitechapel Bell Foundry that cast the Liberty Bell and Big Ben, dating to perhaps 1420.

Cities with clocks and bell towers had a ‘beat’ – a beat of commerce. We could both meet in the market at noon, not some vague idea of mid-day. Efficiency and innovation increased markedly. Cities achieved better economies of scale. As the Worshipful Company of Clockmakers’, founded 1631, motto states, “Tempus Rerum Imperator”, time is the commander of all things.

Smart Ledgers have grown in popularity over the past five years, in large part due to increasing confidence in their performance and security across networks of machines. This confidence is based on their use in cryptocurrencies, yet the ‘jury is out’ on cryptocurrencies, due to economic theory and energy cost objections to their usefulness.

However, commercial users have shown great interest in identity, documentation, and agreement exchange using Smart Ledgers. In these cases, despite the fervour of ‘smart contract’ proponents, so far, the greatest value has been in providing an independent, authoritative record of a document and a time, a timestamp, without giving central third parties too much power.

Central third parties are a well-known approach to handling trade of all kinds. When they work well they earn the sobriquet ‘trusted third parties’. Central third parties typically do three things in financial services: (1) They preserve the definitive set of market transactions. This raises the prospect of charging market participants to ‘get their own data back’. (2) They safeguard the definitive set of market transactions against alteration. This raises the threat of being bribed or rewarded for falsifying transactions. (3) They validate new transactions and authorise their addition to the definitive set of market transactions. This raises the possibility of falsifying assets or admitting corrupt participants.

Further, central third parties frequently become ‘natural monopolies’. A natural monopoly is a supplier whose costs are lower than the alternative of multi-firm provision. Natural monopolies are not inherently ‘evil’, but two aspects are clear. First, a natural monopoly creates at least the three temptations to cheat enumerated above. Before you find this extreme, remember the scale of the FX or Libor scandals just to get started. This is one reason monopolies attract social attention, and in turn regulatory attention. It is also the reason natural monopolists are often paid well by members. If they get caught cheating they put a cushy life at risk.

The second temptation is to extract excessive ‘economic rents’. Economic rents are payments to an owner or factor of production in excess of the costs needed to bring that factor into production. This economics jargon means central third parties can charge much more than things cost. Banks, for example, have long complained about the charges of SWIFT, credit card processors, and exchanges. Switching suppliers in financial services incurs the cost of changing processes for a new supplier, or finding a new supplier with the same level of connectivity, but one of the biggest switching costs is historic data. Often, only the central third party has the authoritative dataset.

The advantage of Smart Ledgers lies not in being cheaper or faster. The advantage of Smart Ledgers is that they allow organisations to work together without giving central third parties a strong natural monopoly. Smart Ledgers do this by giving everybody an immutable copy of the data they need while also reducing ‘switching costs’. To switch to a new supplier, customers need to merely appoint a new central third party, not be hostage to a monopoly on historic data.

And, in many senses, all this is just timestamping.

That is why I welcome this report. While much attention is being paid to the potential applications of smart ledgers, mutual distributed ledgers, distributed ledger technology, blockchains, and the like, a significant portion of the value to date seems to just be providing un-owned timestampers. If so, then it is important that we understand the foundations of time recording and how to integrate ‘time’ properly into our latest commercial ‘beat’, Smart Ledgers.

Professor Michael Mainelli
Executive Chairman, Z/Yen Group
Master, The Worshipful Company of World Traders

World Traders

Executive Summary

A great deal of modern technology is devoted to record-keeping. In datacentres all over the planet, trillions of events of all kinds are being recorded, all logged and timestamped. Different organisations have different standards for the accuracy and precision of their timestamping functionality, some of which are onerous and difficult to meet.

A new technology worth considering in this light is Smart Ledgers – a mutual distributed ledger with an immutable, publicly visible audit trail. Smart Ledgers allow for the storage and future execution of computer programs, so-called smart contracts.

A smart contract is nothing more than dumb code, and code must state its assumptions explicitly. This means that every assumption under which a smart contract runs will be embodied in publicly available code, including regulations connected with timestamping. The issuer of a smart contract cannot afford for its behaviour to be determined by unknown standards. This explicitness benefits all parties: those affected by the contract and the auditor. They can all examine the code at any time.

Throughout the history of timekeeping, two important metrics have been vital to the success of any technology: accuracy and precision. Whenever new timekeeping technologies arrive, their accuracy and precision must be evaluated anew.

Accuracy is the property of a measuring device, which describes how close its measurements are to some fixed reference. Precision, on the other hand, describes granularity of the measurements. In a sense, accuracy is a measure of consistency with external standards, and precision a measure of inner consistency.

The time parameters and timestamps of smart contracts and transactions cannot be compared with one another unless exactly the standards of precision and accuracy are applied to all of them. If a user wishes to post a smart contract to a Smart Ledger, it is vital that the contract abides by the time rules of the ledger. Otherwise, since comparability is impossible, the transactions, contracts and assets cannot exist in the same universe.

In the world of finance, there are stringent regulations that govern the way in which events are timestamped. At the beginning of 2018, the markets in financial instruments directive (MiFID) II regulation came into force in Europe. Among other things, the new rules specify that all transactions are timestamped to a certain level of accuracy and precision. Specifications vary by market. For some areas, the precision and accuracy of the timestamps is on the order of microseconds. In addition, all market participants must record these transaction times according to an approved external clock, so that the timestamps are comparable across different systems.

Later, this report sets out the CUTI requirements for a good timestamping system: it must be Comparable, Universal, Traceable, and Immune. Smart Ledgers, properly used, are capable of meeting all of these requirements.

Smart Ledgers are well-suited to some of these requirements, and less so to others. First, one timestamping-related aspect of MiFID II is complete transparency. The auditor may ask at any time to examine the architecture that a market participant uses when recording timestamps on transactions. This requirement is easily met by Smart Ledgers, which consist of visible code running on distributed hardware. Smart Ledgers are transparent by definition, and easily auditable.

Second, Smart Ledgers are well-suited to the regulatory requirement that all transactions should be stored in such a way that their order can be reconstructed after the fact. A distributed ledger is sequential by definition, and cannot have any ambiguity in this respect. This means that the requirements for highly granular timestamping are met automatically.

The regulations also require that timestamping be fully traceable to a central reference time. The distributed nature of Smart Ledgers makes this a little more difficult, but by no means impossible.

As the standards for timestamping become more and more stringent, we will need to look to new technologies for help. The use of Smart Ledgers may prove a great boon for all market participants. Regulators will find it easy to enforce technological standards, even in areas as challenging as timestamping.

Answer to front cover question: The odd clock out is the one in the bottom right position, which should read 9 o'clock, based on the Cartesian Coordinates of the positions of the clocks.

1. The Use Of Smart Ledgers

Smart Contracts & Dumb Code

A smart contract is a piece of code that can, in principle, embody a set of financial or legal dependencies, and execute some resultant logic in a conditional or scheduled way. Strictly speaking, there is no technical distinction between a smart contract and any other computer program.

The ability to automate the execution of real-world transactions in an intelligent way would result in significant savings in labour and time. In addition, the use of smart contracts could help clarify the workings of the contracts that they model themselves, allowing potential disputes to be resolved ahead of time.

In the financial world, for example, if a complex structured asset is bought and sold, the logic that governs the asset is stored in different ways, in multiple places, for various audiences. Lawyers and investors look at the prospectus that details the behaviour in prose. Quantitative analysts look at equations that embody the prose. They then suggest models to price the asset. Developers design algorithms that implement the models. Traders look at the risks inherent in the model, for hedging purposes. Compliance looks at the buying and selling performed by the traders to hedge the asset. Accountants look at the profit and loss realised by each element of the hedged portfolio.

At every stage in the process, different languages and assumptions are being employed. The asset is written down in different ways and it is up to humans, working in a cross-disciplinary way, to translate from one language to another, hoping that nothing is lost along the way. A single, visible repository for this information would be a clear benefit.

One of the dangers of encoding real-world legal and financial logic in a computer program is that every single aspect of the logic must be explicitly specified in the code. After a contract has been agreed upon and the code put into action, the logic governing the behaviour of the contract cannot be put through a ‘human filter’, in order to test the behaviour for reasonableness. Put simply, smart contracts are made of dumb code.

If any logic is coded shoddily, or left implicit, the smart contract will execute according to its own low-level assumptions, which almost certainly will not match those of an industry practitioner. Any poorly-written code to do with timestamping accuracy and precision could therefore have drastic consequences.

Smart Ledgers

Smart Ledgers are based on a combination of mutual distributed ledgers (aka blockchain: multi-organisational databases with a super audit trail) with embedded programming and sensing, thus permitting semi-intelligent, autonomous transactions. Smart Ledgers are touted as a technology for fair play in a globalised world. 'Blockchain technology' is certainly hyped, and the volatility of related cryptocurrency coins and tokens doesn't help an observer see the underlying robustness of a basic 'internet of record' technology. However, there are numerous projects building trade systems using this technology, with announcements from governments, shipping firms, large IT firms, and the like.

Smart Ledgers allow for not just recording transactions, but also the storage and future execution of computer programs, so-called smart contracts. Smart Ledgers have the potential to transform the ways in which we store and process identities, transactions, debts, and contracts. In a Smart Ledger system, details of posted smart contracts are publicly viewable on the audit trail. The logic of the code that constitutes the contract could be inspected before the contract is bought or sold or otherwise agreed upon. When the contract is posted on the Smart Ledger, it would be obvious whether or not it had been implemented as agreed.

This allows for the examination of assumptions by all interested parties. The actions of the smart contract after it runs are, also, fully auditable. At the same time, the smart contracts are protected by powerful cryptographic methods from being tampered with or altered, after being posted to the ledger.

In a distributed ledger setup, smart contracts are duplicated on every node. The code that embodies the contract runs on all nodes. If the contract detects that some condition has been met and causes a transaction to be executed, this may post another transaction on the ledger – say, the transfer of funds or property of some kind.

In principle, the action taken need not be limited to the ledger on which the smart contract sits. If the MDL has some kind of programmatic access to outside code, then the smart contract can fire any type of transaction. This might include sending a BACS payment in cash, executing a transaction on some other MDL, or buying an asset on a central exchange.

Smart Ledgers & Time

If smart contracts are to be used in real-world applications, they will almost always need some access to a reference time. Most legal or financial agreements are defined or limited by time factors.

Let us say that a smart contract is running a buy order. It is tracking the price of an asset, with an instruction to buy the asset if and when it reaches a certain value. The asset's price is provided by some third party as 'tick' data – i.e. intra-day price information, typically published at frequent but irregular times. If the smart contract is coded to fire when the asset hits the specified price, then knowing the exact time of the published price is crucial. The timestamp on the transaction at the exchange must match the internal time recorded by the smart contract to a high degree of accuracy. Otherwise, it may be that the asset is bought for a wrong or outdated price, and disputes will arise.

Obviously, a smart contract need not be operating in the sphere of financial markets, in which case timestamping requirements may be less stringent. For example, if a smart contract is being used to track the movement of freight,[1] then the times recorded at each stage of the voyage may only need to be accurate to the minute, or even the day.

A salient feature of Smart Ledgers is their super audit trail. Every contract or transaction posted on the ledger is stored in an immutable way, which all participants must agree on. If this audit trail is indexed using accurate, precise timestamping, it will have beneficial implications.

One scenario might be if the Smart Ledger is simply used for historical price data. Historical data is used in the financial industry for assessing trends, calculation of higher-order properties such as volatility and correlation, the calculation of

[1] http://www.longfinance.net/DF/Economic_Impact_Of_Smart_Ledgers_On_World_Trade.pdf

proxy prices for illiquid assets, and the operation of trading strategies (e.g. moving averages). Acquiring and cleaning historical data is often a time-consuming and fiddly process. Having an audit trail of historical prices, all with internally consistent accuracy and precision, would be a significant cost-saver.

Trading With Smart Ledgers

If accurate, precise timestamping is introduced to the world of Smart Ledgers, trading methods could be substantially improved. Currently, an order is routed from the investor to the broker, who attempts to fill the order. The broker can route the order to an exchange or a market maker, fill it internally, or automatically match it using a matching engine. The broker has an obligation to get the best price for their customer. But it is not always clear, after the fact, that this has happened, and it is quite difficult to check. The broker has many orders to fill and will not always fight for every penny on every order.[2]

The presence of a broker causes delays. If an investor makes a decision based on published price data, they would like that decision to be acted on as soon as possible, before the price moves. The broker system may be too slow for investors who are interested in particularly volatile assets.

However, if orders were placed on a Smart Ledger using trusted timestamping, it would be indisputable when an order was placed. Since the method of timestamping and the time of the order are fully visible to everyone, it doesn't matter if the price is applied after the fact! There is no issue with applying the price retrospectively. The bid/offer prices of the asset that prevail at the exact time of the order will be the prices that are applied to that order.

This logic only works when confined to a single Smart Ledger. Attempting to ensure accurate timestamping across ledgers may lead to difficulties. It will be hard to establish precedence between transactions, when each is stamped using potentially a different reference clock, with different accuracy and precision parameters. Reconciling transactions across different ledgers may be impossible. If trading with Smart Ledgers is to become truly universal, then participants should ideally be able to use smart contracts to trade easily across other ledgers.

[2] For a nice breakdown, see https://www.investopedia.com/articles/01/022801.asp

Other Applications

Auctions

A Smart Ledger consists of a chain of events where the current state of the ledger is a result of all previous transactions in that ledger. A classic example of this kind of setup in the real world is an auction, whether this is a classic Sotheby's-type auction, presided over by a man with a gavel, or an electronic auction for advertising space on a website.

A key element of auctions – particularly the old-fashioned type – is that they rely heavily on trust. The auctioneer must trust that the bidders have the money to pay for the goods. Bidders must trust that other bidders have the funds, if the auction is to be meaningful. The bidders also have to trust that the auctioneer has the goods, as well as the right to sell them for the maximum bid. All this makes an auction situation a great candidate for a Smart Ledger, which uses accurate timestamping.

Each bid can be modelled as a smart contract. The contract promises to pay the highest sum it bids, if no further bids are forthcoming. The contract could execute automatically when the auction closes, reducing the need to pursue the winning bidder for payment. If the auction house is holding money in escrow as a guarantee, which is returned to all participants upon completion of the auction, this process could also be triggered with a smart contract.

Fund Management Compliance

The world of fund management is increasingly subject to compliance regulation. Firstly, the levels of risk must remain within the bounds specified in the fund's prospectus. Secondly, any assets explicitly forbidden (for political or ethical reasons) must not be invested in. Thirdly, any investment regulation set out in law must be adhered to. Occasionally rules are breached. Either systems fail to warn traders before they trade or traders ignore the warning. In such circumstances, transactions usually have to be backed out within a certain time. A Smart Ledger would allow investors and auditors to examine positions taken on the portfolio to see when compliance checks failed and whether the backout rules were followed promptly.

In addition to rules covering positions, rules covering transaction timing have become increasingly important in portfolio compliance in recent years. The regulators have begun to take more notice of illegal activities such as wash sales- where a trader sells an asset at a loss for a tax benefit and then repurchases it soon after - and cross trades - where a broker matches buy and sell orders for two clients without recording either transaction on the exchange.

Since these phenomena are defined by time limits, accurate timestamping on a publicly visible Smart Ledger would help with quick, automatic detection.

Smart Ledgers & The Law

As Smart Ledgers become more common in the financial world, they will be scrutinised by regulators very closely. So far there has been no attempt at writing financial regulation that explicitly covers Smart Ledgers. There are still terminological issues to be addressed before this can happen. A few sovereign entities have recognised the legal status of smart contracts – the state of Tennessee being a recent example[3] – but the uptake has not been fast. The concept of a 'contract' has a special meaning in law, and it is not clear that all the shades of meaning can even be encapsulated in code, or whether computer scientists and lawyers are talking about the same thing.

Before Smart Ledgers can be taken seriously, regulatory bodies will need to be satisfied that Smart Ledgers can function in strict accordance with existent financial regulation. There exists a large body of new regulation that deals in transaction timestamping. Recent regulation, such as Europe's markets in financial instruments directive (MiFID) II, explicitly details compulsory minimum standards for timestamping, which will be an important factor in assessing any new financial technology.

It is, therefore, vital to think about whether Smart Ledgers can meet new requirements.

[3] https://www.coindesk.com/blockchain-bill-becomes-law-tennessee/

2. A Brief History Of Timekeeping

On The Hour

As long as there has been civilisation, there has been the marking of time. The earliest astronomers, in the Fertile Crescent and in China, used the constellations to divide the night into equal chunks of time. Their zodiacal systems were similar, and this base-twelve division of space yielded a base-twelve division of time. The Ancient Egyptians were the first to divide the hours of darkness into twelve equal units, based on the positions of constellations. By the first century B.C., Greek astronomers were accurately dividing the day/night cycle into 24 equal units, using a combination of water clocks and sundials.

In China, the use of water clocks was widespread, and by the time of the Tang dynasty (600 A.D.), they, too, were dividing the whole solar day into 24 units. As well as its astrological properties, the number twelve is ideal for fractional measurement, as it divides neatly into halves, thirds and quarters.

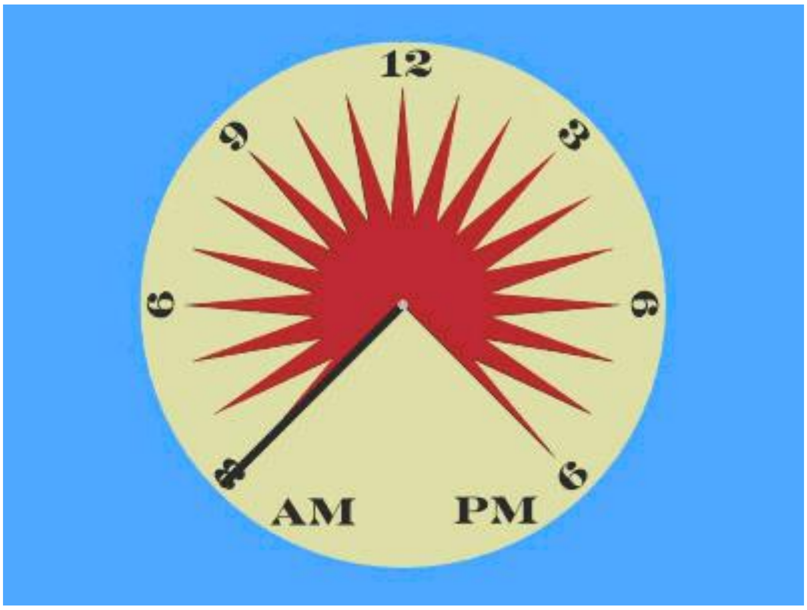

Figure 1: A simple sundial[4]

By Roman times, the time of day had become firmly embedded in public life. The market bell marked the time, every 3 hours roughly from 6 am to sunset. In

[4] https://commons.wikimedia.org/wiki/File:Equatorial_sundial_topview.gif

medieval Europe, the Church used the Roman divisions to specify times of prayer, such as matins, lauds, and vespers. These ‘canonical hours’ were now marked by church bells. The use of sundials allowed (at least on sunny days) for the offices of the church to be kept accurately.

By the 17th century, pendulum clocks had been introduced, and soon after that the balance spring was applied to pocket watches. This greatly improved their accuracy. The hour had been precisely divided into minutes and seconds. Accurate and portable timekeeping was within the reach of the layman.

On The Ball

The implications of portable time were immense. The question of longitude had long plagued navigators. Measuring a ship’s east-west position at sea depended on being able to compare the time of the observed sunrise with the known time of sunrise in Greenwich. A standard pendulum-driven clock was useless since the pitching and rolling of the ship upset the swing of the pendulum. With the advent of compensatory mechanisms and spring-driven clocks, ships could carry their clocks far from home, secure in the knowledge that they would not lose (much) time.

In order that visiting ships would be able to synchronise their clocks, many harbours offered the service of a ‘time ball’. This was a large ball on a pole, mounted on top of a tall building visible from the water. Every day at exactly 1 pm, the ball would slide from the top to the bottom of the pole.

The use of a ball to show the time of day dates back to antiquity. A modern descendant of it is seen at midnight on New Year’s Eve in New York, when a lit ball is lowered down a pole on top of One Times Square.

As British maritime influence grew, the use of Greenwich time as a reference became widely adopted. A ship’s time difference from Greenwich reflected its spatial distance from Greenwich. The distance east or west of Greenwich, therefore, became the standard worldwide measure of longitude. This was formalised in 1884, when an international committee decided to use the line of longitude through Greenwich as the Prime Meridian.[5]

On an everyday level, the ability for many people to carry their own watches meant that they were no longer dependent on the church clock of the local parish. This opened up a great deal of commercial possibilities, including financial agreements that depended on time, as well as firmly establishing the idea of an hourly wage. The definition of a mutually agreed time has always had an intimate relationship with finance. Time may not be equal to money, but it is at least proportional.

Over The Wire

Over time was portable, it was in a sense, no longer local. In the Middle Ages, the lack of accurate portable timepieces meant that a traveller had no way of knowing whether the parish clock at one location was synchronised with one at another location. Soon after the invention of the electrical telegraph, Gauss experimented with using it to synchronise clocks in different locations.

This idea really took off when commercial rail began. The telegraph and the railways grew hand-in-hand. As the rail networks grew, first in Britain, then elsewhere, the use of the telegraph exploded. At first, the telegraph was used to send messages up and down the lines to points operators, station masters, teams manning repair works, and so on. But as the timetables became more complicated, the question of standardised time became more of an issue.

The public clocks in towns and villages in Britain were periodically reset using sundials. Thanks to differences in longitude, even within England, times on geographically separated clocks could be up to 30 minutes apart. In the United States, which spans multiple time zones, the differences could be up to an hour. In order to ensure punctual train services, not to mention avoiding accidents, it became a matter of urgency to synchronise the clocks across the network.[6]

[5] The French abstained, although by this stage even they were publishing nautical almanacs that measured longitude from Greenwich, rather than Paris.

[6] This occasionally caused friction when the residents of a far-flung town refused to adjust their clocks to Greenwich time. Both Wordsworth and Dickens expressed unease that an artificially imposed central time had superseded the more natural measure given by the sun, purely for the convenience of railway operators.

This was not as simple as using a sundial at Greenwich to set the time. Thanks to the Earth’s elliptical orbit, the time at which the sun is highest in the sky can vary by as much as quarter of an hour either side of the mean. The notion of ‘mean time’ ensures that 12 noon is set to the average time in the day when the sun is overhead, over the course of the year. By 1855 a time signal was being sent from Greenwich by telegraph throughout the whole railway network to allow for synchronisation to Greenwich Mean Time (GMT).

Railways in the United States and on the Continent soon followed suit, standardising their national times (often to Greenwich), and synchronising their networks.

3. Timekeeping Standards

International Atomic Time

International Atomic Time (TAI, for Temps Atomique Internationale) is the international community's principal measure of time. More specifically, it is a measure of 'proper time', which means that it measures the actual passing of seconds, rather than referring to the movement of the Earth or the Sun.

The scheme is administered by the International Bureau of Weights and Measures (or Bureau Internationale des Poids et Mesures). The time calculated is typically accurate to within 10^{-15} seconds, or one femtosecond. The measurement is generated from the average of over 400 atomic clocks globally. The most common type of atomic clock is the caesium clock, but others are also used. The International System of Units (SI, abbreviated from the French Système International (d'unités)) definition of 'a second' is based on the nature of the radiation produced by the caesium-133 atom, hence the popularity of caesium atomic clocks.

TAI time is not published live. Instead, a monthly publication called "Circular T" is made available. This lists the amounts by which each contributor to the average deviates from that average. This strategy addresses the fact that there is uncertainty in the offset time of each contributor's clock. Offset of any clock is subject to drift, changing the offset value over time. Regular recalculation and re-publication of Circular T ensure that this is corrected.

Using Circular T, the live time published by any of the contributing atomic clocks can be adjusted by subscribers in order to match TAI.

Coordinated Universal Time

Many observatories and other institutions around the world participate in the TAI scheme described above, broadcasting on public radio their measurements of time under an allotted UTC code.[7] The UK's National Physical Laboratory uses

[7] The acronym UTC, used globally, does not actually stand for anything. It is a combination of TUC (*Temps Universel Coordoné*), and CUT (Coordinated Universal Time).

the code UTC(NPL), whereas, in the States, the National Institute of Standards and Technology uses UTC(NIST).

However, these times are not each participant's estimate of International Atomic Time, but a measure of Coordinated Universal Time, or UTC. This time standard is running, at present, 37 seconds behind TAI. The difference between the two standards is that while TAI is a measure of 'proper time', UTC retains a link to solar time. The UTC standard ensures that midday is, on average, the time of day at which the sun is due South (in the Northern Hemisphere). The two times differ because the rate of the Earth's rotation is irregular. Therefore, the length of the solar day is almost never exactly the same as the ideal 24 hours, as measured by an atomic clock.

To adjust for this, the leap second is used to correct UTC. In 1972, UTC and TAI differed by exactly 10 seconds (for various reasons). Since then, a leap second has been added 27 times, so that UTC is now 37 seconds behind TAI. In principle, leap seconds can be added or removed. Thus far, the Earth's rotation has only slowed down, but it may speed up at some point.

When leap seconds are added, it tends to be at midnight UTC time. An extra second is added at the very end of the day. On those days, the clock reads 23:59:59, then 23:59:60, then 00:00:00. In a global high-frequency financial system, it is vital that a central clock is used to set this time, that the mechanism is understood, and that the central reference time adjustment is propagated to dependent clocks. The last time this occurred was on 31 December 2016.

GPS Time

GPS Time is yet another time standard, used by the Global Positioning System (GPS, discussed below).

The clocks on board GPS satellites and in ground stations are set to the GPS standard. This is not affected by leap seconds, and therefore tracks TAI. However, GPS time was synchronised to UTC on the 6th of January 1980. It is therefore ahead of UTC by 18 seconds, and behind TAI by 19 seconds. If the GPS receiver is to be used as a clock, this must be taken into account.

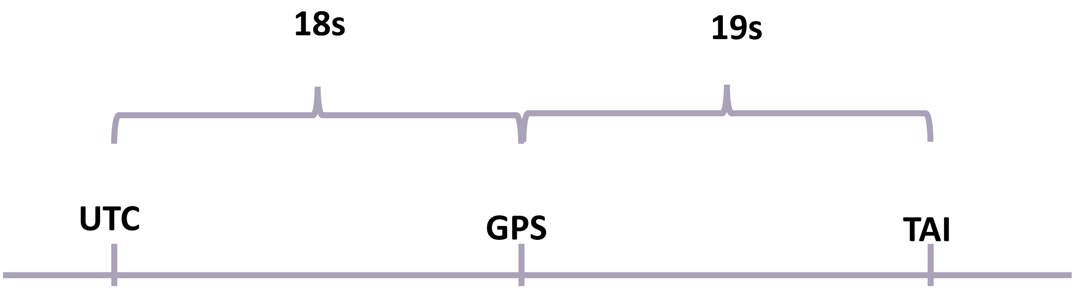

Figure 2: UTC, GPS time, and TAI shown on a timeline

4. Broadcast Time

National and international timekeeping standards are useless if they are not made easily available to institutions who wish to synchronise with them, as well as with one another. Central reference times are therefore publicly broadcast using different schemes.

It is important to distinguish between accuracy of a time signal and its precision. A useful analogy is that of a dartboard. The bullseye of a dartboard is the true value that we would like to measure as closely as possible. If a group of darts are loosely clustered around the bullseye, they have accuracy but low precision. If they are tightly clustered in some other area of the board, they have precision but low accuracy.

One consequence of this definition is that precision has nothing to do with the true measurement of time. Precision is a measure of the granularity of the data we are presented with – the number of decimal places, for example. A watch might measure time with the precision of 1 second, but it is usually inaccurate by a couple of minutes.

The 'time ball' discussed above also illustrates this. Say the ball is raised at 1 a.m. and dropped at 1 p.m. This may be accurate to the second (or even better), but since a movement of the ball only happens twice a day, you cannot "read the time" from the ball with better precision than 12 hours.

'Time From NPL' Signal

The national reference time for the UK has the code UTC(NPL). It is broadcast from a radio transmitter in Anthorn, Cumbria. It is kept accurate using three atomic clocks on-site, and takes its reference from the TAI standard, as described above. The signal can be received all over the UK, and is also widely used in parts of western Europe. A clock with a radio receiver set up to synchronise with this signal will always display the correct UK time, and will automatically adjust for leap seconds and, optionally, Daylight Savings Time.

The signal has a very low bandwidth. Only two bits are transmitted per second. At the beginning of each minute a marker is broadcast, in the form of a long interval. Over the course of the subsequent minute, the two bits per second transmit the full date, day of the week, and the time in hour-and-minute form.

The signal, also, carries information about the time's offset with astronomical clocks – i.e. how close we are to a leap second being necessary.

The signal is accurate to within 1 millisecond of UTC, although it only has a precision of one second.

Global Positioning System

The Global Positioning System (GPS) is the most widely used system for both positional navigation and timekeeping. This is a satellite system, owned by the US government, which can be used by any GPS receiver to calculate location information and time. It became fully available for civilian use in 2000.[8]

Using the satellite signals, a GPS receiver can calculate its position (three values – latitude, longitude, and altitude) and the time error of the receiver clock (one value). Since there are four values to be found, the receiver needs to have an uninterrupted line of sight to at least four satellites.

Each satellite periodically broadcasts the three values representing its own position, as well as the time measured by its onboard atomic clock. These are highly stable clocks, resynchronised with each other and with ground clocks every day. The time broadcast is the GPS time standard, discussed above, currently ahead of UTC by 18 seconds.

Using the times and positions reported by four satellites, and the known speed of light, a GPS receiver can calculate its distance from each satellite, and the time error of its own onboard clock.[9]

It does this by solving four simultaneous equations with four unknowns, we want to find t_0, x_0, y_0, z_0; the time and positional values of the GPS receiver. The four equations we use to solve for these values can be expressed as:

$$c^2(t_0 - t_n) = (x_0 - x_n)^2 + (y_0 - y_n)^2 + (z_0 - z_n)^2 \quad \text{for n = \{1,2,3,4\},}$$

[8] https://www.gps.gov/technical/ps/2008-SPS-performance-standard.pdf

[9] In practice, a receiver can "see" an average of nine satellites from any given point on the Earth's surface. This redundancy allows for error-checking.

where t_n , x_n , y_n , z_n are the times and positional values of four GPS satellites``````````````, for n = {1,2,3,4}, and *c* being the speed of light.

As encoded by those four equations, the variables representing location and time are inseparable and interdependent. If there is a time delay in the signal transmission, this will result in a proportional shift of the receiver's calculated position in one or more dimensions. Similarly, any positional error coming from the satellites will result in the timing signal being misinterpreted.

Even if there are no errors in the signal, the accuracy and precision of the time inferred by GPS receivers are directly related to the resolution of the position measured by the satellites. The positional granularity of the system is to within a few metres (on average, about 15 metres). This corresponds to a time precision of tens of nanoseconds. The distribution of the GPS time signal is guaranteed to have a standard deviation of 97 nanoseconds from UTC. Propagation and receiver errors mean that this range moves up to within 200 nanoseconds of UTC – not forgetting the 18 second adjustment for GPS time.[10] Indeed, the calculation by the satellites of their own position is dependent on a system of meticulously averaged atomic clocks both on Earth and in orbit. Timekeeping and time comparison are at the core of the whole system.

Time Synchronisation

After an organisation has set up a reliable link to a verified source of time, the next issue is to synchronise its internal systems to that time.

This problem is not unique to computer networks. Producers of live TV have long been aware of the difficulty of synchronising video with other video for the purposes of smooth live editing, as well as synchronising video with audio. A technique known as Genlock has emerged, where one video device generates a signal that other devices must 'lock' to.[11]

[10] https://www.gps.gov/systems/gps/performance/accuracy/

[11] https://en.wikipedia.org/wiki/Genlock

On Local Area Networks (LANs) and on the Internet, the challenge is to synchronise every machine on the network to within a certain tolerance of the approved reference time. A number of protocols and technologies have been developed to address this. Many of these techniques make use of 'smart clocks' – a clock that compensates for offset, frequency error, rate of change of frequency error. The provision of these technologies is an industry in itself.

Network Time Protocol (NTP)

Dating back to 1985, NTP is still in current use in some areas. The timekeeping network is structured as a hierarchy, in levels of accuracy called 'strata', as shown below.

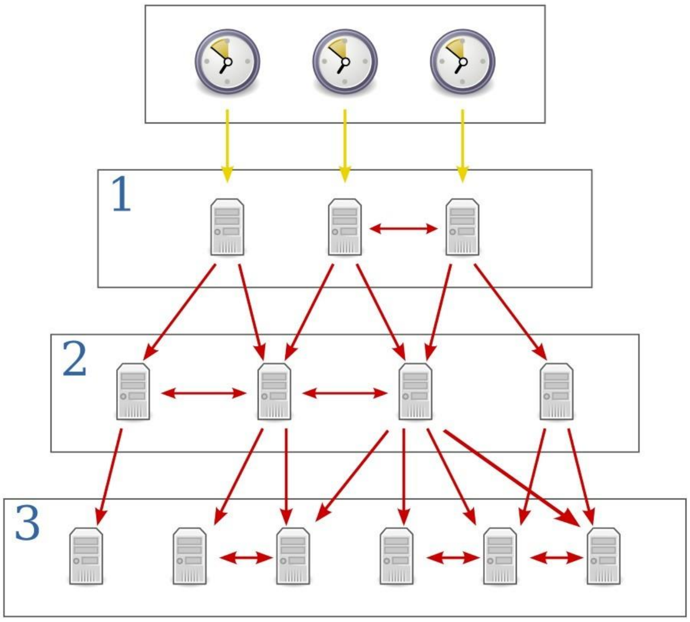

Figure 3: The Network Time Protocol[12]

At the top, Stratum 0 consists of reference clocks – atomic, GPS, or radio clocks. Below these on the hierarchy is Stratum 1, consisting of what is known as the primary servers. Each one is synchronised to a reference clock, as well as being linked peer-to-peer as a sanity check.

Strata 2 and 3 are linked in a similar subsidiary tree-like structure to the stratum directly above each in the hierarchy, and also linked peer-to-peer. In principle,

[12] https://en.wikipedia.org/wiki/Network_Time_Protocol

the number of Strata can rise as high as 15, but in practice they rarely go higher than 4.

NTP's hierarchical structure ensures that the reference signal can reach every node on the network very quickly. Each client uses a polling mechanism to calculate the latency between it and a chosen server. First a request is sent from the client, with a time stamped on it. When the server receives this, it records the time the packet is received. It then sends a timestamped response back to the client. The time of receipt is recorded by the client.

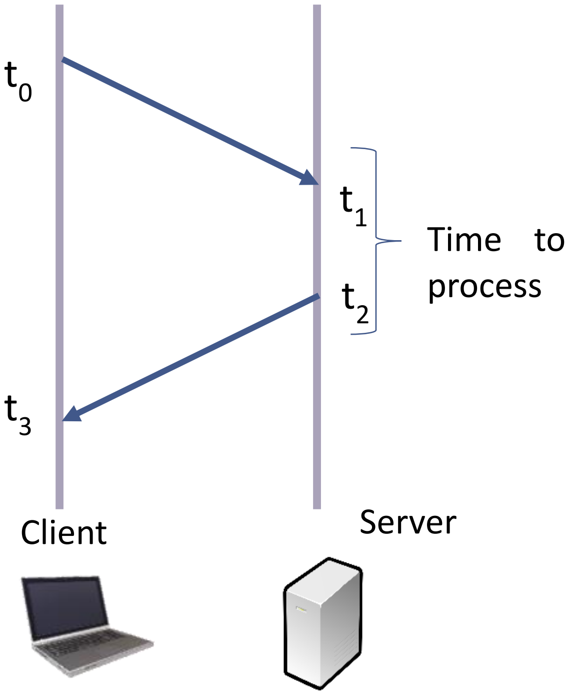

Figure 4: A client polling a server to calculate the offset and latency

In order to check synchronisation between client and server, two values must be calculated: the offset, which is the difference between the time recorded on the client and the server; and the delay, which is the total real time taken for a message to travel from the client to the server and back, ignoring processing time.

If:

t_0 = the time the packet request was sent from the client,

t_1 = the time the packet request was received by the server,

t_2 = the time the packet response was sent from the server, and

t_3 = the time the packet response was received by the client,

then the offset, θ, is computed as follows:

$$\theta = \frac{(t_1 - t_0) + (t_2 - t_3)}{2},$$

and the delay, δ, is computed as follows:

$$\delta = (t_3 - t_0) - (t_2 - t_1).$$

Each client in the hierarchy polls 3 or more servers, and uses the timestamps returned to compute the clock offset and the round-trip delay when sending a message to each server.

Once this is done, each client determines which time sources it should use for estimating the correct time. The time data received from all servers is subjected to a statistical method called the intersection algorithm. Using this, the client calculates the most likely estimate of the correct time, and what the confidence interval for that estimate is.

The algorithm does this by systematically excluding 'false tickers' – times that are outliers from the average – in such a way that the confidence span for the time is minimised for the maximum number of noisy sources.

NTP timestamps are 64 bits long: 32 bits for counting seconds, and a further 32 bits for fractional seconds. It therefore rolls over every 2^{32} seconds (136 years) and has a precision of 2^{-32} seconds. The epoch – the point from which the time is measured – is the 1st of January, 1900.

The accuracy of the NTP system is in the millisecond range in an internal LAN, although this degrades to tens of milliseconds over the Internet. This accuracy is widely felt to be insufficient for modern timestamping applications, particularly in the financial world. In addition, in recent years, the protocol has been criticised for being insecure. If a client in the hierarchy were compromised, the transfer of timestamps might be corrupted. NTP could be made more secure by using cryptographic methods to sign each packet, but the computational overhead required would unacceptably affect the performance. The NTP method has therefore been used less in recent years.

Precision Time Protocol (PTP)

The current best industry standard, published under IEEE-1588, is the Precision Time Protocol, or PTP. Using this method, a network can be synchronised to nanosecond levels of accuracy.[13]

This is achieved by using a sophisticated algorithm to calculate latency. We saw that under NTP delay was simply defined as a round-trip time between client and server. However, this is not necessarily a constant or predictable number. Packets sent across a network are subject to all sorts of discontinuous delays in the form of queues. The network hardware is constantly routing packets everywhere, and this means that the time taken for a given packet does not necessarily have a linear relationship with the distance it has to travel over the network. Nor does a packet travelling between A and B have to take the same route in reverse as a packet sent from B to A. To use round-trip delay to estimate latency is, therefore, not sufficiently precise.

To add more insight to the method of calculation, PTP therefore uses hardware timestamping. In order to synchronise a network, PTP requires that the switches and routers on the network are PTP-enabled. This means that they have on board transparent clocks. When a timestamped PTP message passes through a transparent clock, it updates the timestamp on that message, to adjust for the time the message spends in the switch or router. All the unpredictable delays caused by switches and routers are removed. Once the queue-based latencies have been removed, the timestamp on the PTP message is a genuine reflection of the time the message spent traversing the network. This means that the estimate of latency between two network nodes is much more accurate.

[13] A prominent vendor at present is Meinberg. See their take on the IEEE-1588 at http://blog.meinbergglobal.com/2013/09/14/ieee-1588-accurate/

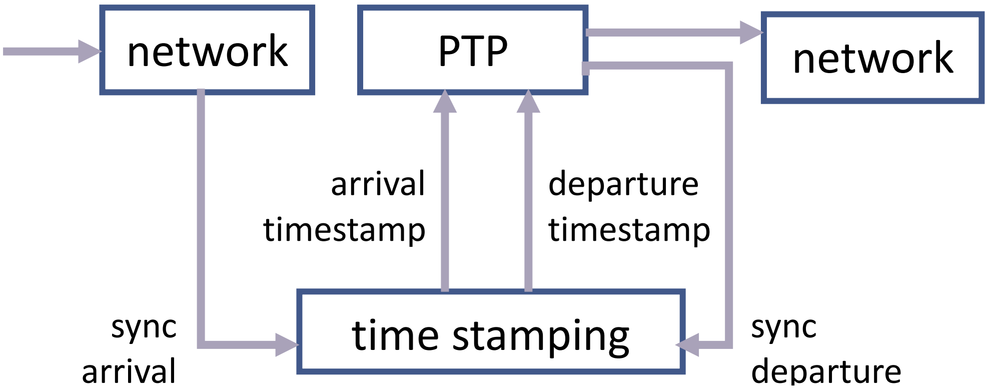

Figure 5: A transparent clock system

Using this measure, PTP employs an algorithm called the best master-clock (BMC) routine to select a grandmaster reference clock. This is in contrast to NTP, where the network is more rigidly hierarchical. The ability of the network to select its own reference makes the system more accurate. It also means it is more robust since, if the best clock goes down, the next best can be used.

NPLTime®

NPLTime® is a service offered by the UK's National Physical Laboratory (NPL). Unlike the radio signal described as 'Time From NPL' (above), which is a public service, NPLTime® is a paid-for service marketed directly at the financial sector. It is specifically advertised as being MiFID II-compliant.

The NPLTime® service consists of a secure fibre optic link direct to the central NPL atomic clocks. The protocol used is PTP. In order to use the NPLTime® service, therefore, it is necessary for an institution to have high-spec optical switches and routers that are PTP-enabled.

NPL certifies that the time signal is accurate and traceable to UTC. The service is capable of synchronisation to UTC(NPL) to an accuracy of 100 nanoseconds.[14] More than that, NPLTime® certifies the time at the user's location. This means that the user doesn't have to worry about latency effects outside their own network.

[14] http://www.npl.co.uk/upload/pdf/npl-ubs-white-paper.pdf

How the time source is used is down to the user. NPL advise customers to put their master server in a rack next to the NPL receiver, then synchronise the network. Beyond that, the use of the time signal is not part of NPL's remit. The time is guaranteed accurate, but there are no restrictions on how the time is used, and no guarantee that the timestamps purportedly from NPL have not been interfered with after the fact.

This service completely eliminates reliance on GPS time. The GPS system is free and universally available, and as a result has become the backbone of the whole financial system. Recently, concerns have been raised about its vulnerability to attack[15], so the ability to remove this risk is valuable.

The Future

As computer hardware gets faster, and networks become larger, the time synchronisation problem becomes more and more challenging. In high-frequency trading environments, the levels of accuracy and precision required will only become more demanding. If two timestamps that are not using synchronized clocks are to be compared, large financial losses can accrue.

Regulation will have to keep up with technology as it develops. It cannot be done ahead of time. When the MiFID II regulation (discussed below) was being drafted, the regulators attempted to future-proof the rules. It was, initially, proposed that timestamping be done with nanosecond-level precision, but the regulators were forced to withdraw this in the face of protest. Even the levels of deviation of UTC time providers do not match TAI to within those limits.

The Financial Information eXchange messaging protocol (FIX) is the technology which underpins most of the world's electronic trading communication. The message format currently supports picosecond (10^{-12}seconds) accuracy. Since the accuracy of services like NPLTime® is only advertised as accurate down to the nanosecond, presumably picosecond-level accuracy will be sufficient for FIX's needs for the foreseeable future. FIX have a clock sync working group, that looks closely at these questions.

[15] See https://qz.com/1106064/the-entire-global-financial-system-depends-on-gps-and-its-shockingly-vulnerable-to-attack/

5. Secure Timestamping

It is sometimes the case that a document or transaction needs timestamping in such a way that it can be shown that the time has not been manipulated. This manipulation could happen at the time of stamping, where the wrong time might be attached to the document, or afterwards, where the timestamp might be changed by some malicious party.

When considering solutions, it is important to remember that they should not need a large overhead. Smaller institutions may need a secure timestamping mechanism that does not require them to implement an advanced synchronisation or time subscription scheme. Ideally, timestamping should be a cheap, secure, verifiable public resource.

Centralised

The X9.95 standard from ANSI[16] specifies a protocol for the secure use of timestamps. It relies on the existence of trusted Time Stamping Authorities (TSAs). These are third party entities who are recognised as providers of impartial, reliable timestamps.

In order to timestamp some content, the following is done:

- The user creates a hash of the content. This is a string of text that uniquely represents the content, but the content cannot be calculated from it.
- The user submits their ID and the generated hash to the TSA.
- The TSA generates a timestamp.
- The TSA generates a token, consisting of the submitted hash and the timestamp, and a digital signature for these two pieces of data.
- The token is sent back to the user, who stores it with the original document.

At some later time, if a verifier wants to verify a timestamp, the following is done.

- The verifier compares the digital signature of the token against the TSA's published certificate. If they match, the verifier knows that the token was issued by the TSA.

[16] https://webstore.ansi.org/RecordDetail.aspx?sku=ANSI+X9.95-2016

- The verifier independently hashes the content and compares the hash in the token with the hash of the document. If they match, then the verifier knows that the document hash was used to generate that token.
- The verifier reads the timestamp contained in the token. They know that the timestamp and the hash were used in the generation of the digital signature, which means that the TSA did indeed receive the content at the given time.

An example of a centralised timestamping scheme using Smart Ledgers is the MetroGnomo service provided by the States of Alderney.[17] MetroGnomo is a central 'transmitter' that broadcasts to 'receivers' that record the time. Decentralised timestamping schemes using Smart Ledgers are also possible.

Linked Centralised

A problem with centralised schemes is that users must have absolute faith in the TSA. If the TSA's timestamps become corrupted, every single document that is stamped is in doubt.

One way of mitigating this risk is the use of linked timestamps.[18] Recall that when client n sends a hash of the document for stamping, they also submit their identification (ID_n) to the TSA. In a linking scheme, the TSA not only issues the token to a client, but, also sends the identification of the previous requester (ID_{n-1}). Then, when the next client (n + 1) requests a token, the TSA sends their identification (ID_{n+1}) back to the original client n.

This means that every client has a link to the previous client and the following one, creating a chain. Given sufficient volume of requests, this makes it difficult for the TSA to spoof the timestamps, since the ordering of stamps can be reconstructed by traversing the chain.

[17] https://www.metrognomo.com/

[18] See https://www.esat.kuleuven.be/cosic/publications/article-122.ps for an overview of timestamping solutions, including a linked scheme.

Distributed

Linking schemes improve the reliability hugely, but the TSA remains a single point of failure. In recent years, work has been done on the concept of distributed timestamping. A number of schemes have been proposed that use a consensus model very similar to that of distributed ledgers.[19]
Distributed schemes work as follows:

- Clients submit hashed documents to a random selection of the available servers.
- Each server receives a number of requests from different clients over a certain time period, known as a 'round'.
- The servers broadcast all the hashed requests to each other, so that every server is aware of every request submitted by every client during the round.
- The concatenation of all the requests, and the timestamp of the round, can be used to make a global timestamp representing that round. Each server calculates this. The agreed value is the published global timestamp of that round.
- Each client is issued a token by the servers that received the request from that client. This token consists of the global timestamp, the submitted hash, the time, and a client-specific timestamp.
- If, due to latency, two servers receive the same client's requests during different rounds, the consensus model ensures that the first request will be honoured.
- The relationship between the array of all requests and the client's request is such that when combined and hashed in a specific mathematical way, the result equals the published global timestamp.
- If a verifier wants to ensure that a client's document was timestamped during a certain round, they can perform this hashing operation, and then check that the result is as expected.

Distributed schemes of this kind are very difficult to manipulate, since they operate on a consensus basis. However, they are limited in the precision of the

[19] See https://bravenewcoin.com/assets/Whitepapers/Improving-Time-Stamping-Schemes-A-Distributed-Point-of-View.pdf for a distributed proposal.

timestamps they can provide. The timestamp cannot be any more granular than the length of each round.

6. Desirable Properties Of A Timestamping Scheme - CUTI

It is clear that for a timestamping system to pass muster, it has to be robust enough to measure up to the practical requirements of existing applications, as well as the rules and regulations. This is true of any new solution, whether it is a brand-new idea like Smart Ledgers, or simply a rearranged version of old infrastructure.

Based on some consideration of existing timestamping schemes, regulation and common sense, we have come up with the following neutral properties which we believe a useful timestamping system should have. Call them the CUTI properties. Timestamping schemes should be:

- Comparable
- Universal
- Traceable
- Immune

Comparable

Comparability is fundamental to timekeeping of all kinds. Ultimately, a single time is useless on its own. Timestamps only become useful when being compared to each other. The time of an event, once recorded, will at some point be read either by a human or a computer, in order to compare it to the times of other events. Either an asset is doing some action on a schedule, or it is doing something depends on some detected event, or it is waiting for a certain 'maturity' time. In almost all cases, in order for a contract to execute, there will be a need to compare two times. A timestamping system that records the times of events will not only have to use a universal clock to guarantee *accuracy*, but, provide a specified level of precision for each timestamp. Without these guarantees, there is the danger of disputes when attempting to compare two times.

Universal

A timestamp on a document, smart contract or transaction must be independent of the technology that applied it. Any two systems that interact are going to have to be able to marry up their timestamps with each other, no irrespective of technological mechanisms used by each one. A central, frequently synchronised universal clock will need to be used. The regularity of synchronisation of each network, as well as the infrastructure used to perform it, may also need to be agreed.

Traceable

An important aspect of any timestamping service is guaranteed traceability. It is required by MiFID II that a market participant can demonstrate that time on their clock is traceable to UTC.[20]

This is an onerous requirement. To ensure time traceability, every step back to the reference clock must be known and documented. There must be a full beginning to end view, with the potential error at every point known and documented.

It is easier to guarantee traceability with services such as NPLTime®, which certifies the time signal's traceability to UTC.[21] To ensure traceability with GPS is more difficult. To do this, one needs to calibrate the receiver to know the offset, and then continuously monitor the receiver.

One element that can break traceability is latency. In order to ensure traceability to a central atomic time source, a user will need to have an appreciation of the latencies present in its own internal and external systems.

Immune

The value of time is the ultimate neutral reference point for many applications. A large infrastructure has been set up worldwide so that any system plugged into it can identify what the time is. It is therefore incumbent upon anyone recording the time of an event to ensure that the timestamp they use is immune from interference by a third party. A timestamp should be impossible to change after it has been applied.

[20] https://www.esma.europa.eu/file/16989/download?token=qCLY-NSq

[21] http://www.npl.co.uk/upload/pdf/time-traceability-for-the-finance-sector-factsheet.pdf

Smart Ledgers & CUTI

It is worth examining Smart Ledgers in light of these properties. Will the technology be able to live up to the standards that both regulators and users expect of existing centralised technology?

Table 1: The CUTI properties necessary for a reliable timestamping scheme.

Property	Comments	Score
Comparable	This depends on two major factors – precision and accuracy. The comparability of the time parameters of a smart contract, either with clock time, price data or other timestamps will be limited by the precision and accuracy available to other systems which the smart contract must interact with. This is true of any technical implementation of time logic, but Smart Ledger precision is also affected by the latency inherent in distributed technology.	*Average*
Universal	As long as the users of Smart Ledgers agree on timestamp conventions and authorities, timestamps can be used in all contexts. Cross-ledger comparison remains a problem, although perhaps no more serious than comparisons across market are at the moment.	*Good*
Traceable	If a customer uses the distributed ledger to perform the timestamping, then every node on the system will have to be guaranteed traceable - for example by having each node synchronised with a system like NPLTime®. However, if a distributed scheme is used to aggregate timestamps and produce a collective number, it will have to be considered whether this aggregation breaks the traceability guarantee. This may be far from obvious!	*Good*
Immune	Thanks to the super audit trail and consensus model, as well as the crypto technology used by Smart Ledgers, nobody should be able to change a timestamp after it has been submitted to the chain.	*Very Good*

Test Case – NPL, TMX, Z/Yen, Strathclyde University[22]

In August 2017, it was announced that a group of researchers from the National Physical Laboratory (NPL), the Toronto Stock Exchange (TMX), Z/Yen Group, and Strathclyde University had successfully timestamped stock trades and recorded them directly on a distributed ledger.

The times were in UTC, synchronised with NPL's atomic clocks, and the project managed to record over 20 million transactions from three hours of trading to the ChainZy smart ledger system (note – provided by Z/Yen). The team created a timestamping engine using ChainZy's woven-broadcasting architecture to test the recording of high-frequency trading transactions.

The test used nanosecond resolution high-frequency data from the TMX located in Interxions's London Data Centre with support from Hyperneph. The researchers timestamped smart contracts of various lengths, written to execute a series of buy and sell instructions. These were either logged with NPLTime®, using the atomic clocks at NPL, or logged with UTC plus a randomly generated time lag. The orders were then sent to a central clearinghouse also operating on UTC and written onto a ChainZy distributed ledger.

The ChainZy architecture had a capacity for 25 billion transactions per day on the test rig used. In fact, the research team estimated that the highly scalable architecture was capable of perhaps 1 trillion transactions per day if deployed at levels comparable to popular cryptocurrency systems. The speed of the system is a great riposte to the idea that distributed ledgers have an unacceptably high latency.

The research also discussed issues of transaction processing. A challenge for distributed systems is to process transactions in the correct order. This is necessary to avoid 'front-running', where a market participant uses advance knowledge of future liquidity, in order to make a profit.

The research found that while precision timestamping was required by the ledger to record the order of each transaction's arrival, it was not sufficient to do this on the ledger only. To prevent front-running, transactions must be ordered for execution by booking time, not by time of receipt. This research is the subject of a journal article by Strathclyde University.

[22] Basu D, Broby D, Arulselvan A. The role of precision timing in stock market price discovery when trading through distributed ledgers. Working Paper, Glasgow: University of Strathclyde. 2017 Jan 23, p. 1-23. https://pure.strath.ac.uk/portal/en/publications/the-role-of-precision-timing-in-stock-market-price-discovery-when-trading-through-distributed-ledgers(7af1860f-c758-4699-94ad-993b5025a364).html

7. Smart Ledgers & Timestamping

A. Posting Smart Contracts

When posting a smart contract on a distributed ledger, it is fundamental that the smart contract is accurately timestamped at the time of posting. The distributed nature of the Smart Ledger architecture poses special problems for accurate and precise timestamping of smart contracts.

If a network is based on a ‘hub-and-spoke’ or ‘star’ model, with all nodes posting transactions to a central server, then we can say that an event is deemed to have occurred when the server receives it. The server can then broadcast that event to the rest of the network.

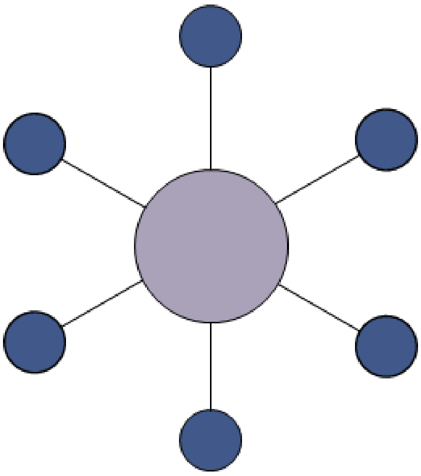

Figure 6: A hub-and-spoke network

However, a distributed ledger consists of a network of nodes. There is no central authority. The audit trail is reached by consensus. In this decentralised ‘mesh’ model, a node broadcasts a transaction to every other node in the network, via as many intervening nodes as necessary.

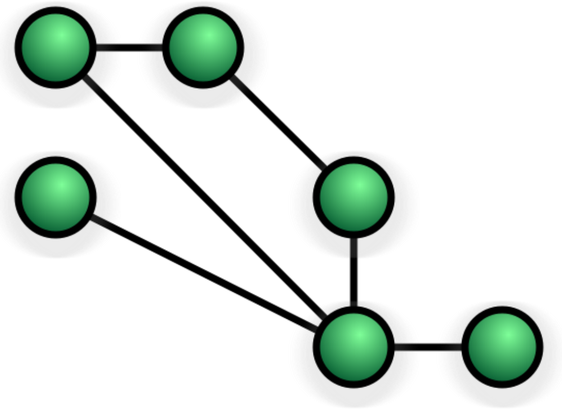

Figure 7: A distributed mesh-style network[23]

The distributed nature of the network means that for a given node, all the other nodes in the network are different distances away. This means that the signal to perform a transaction reaches the other nodes with differing latencies. A client's broadcast of a transaction will eventually reach all nodes on the network, but each node will receive the information at different times. If a blockchain-style architecture is used, it may take some time before the nodes reach consensus on the transaction and store it immutably. There is therefore no single time at which a Smart Ledger transaction becomes official. We cannot use the time the transaction is 'received' as the timestamp, since the transaction undergoes a process of distribution and duplication before being saved down.

An obvious solution is for a client to timestamp its own transaction, then to post the timestamp and the transaction details to the ledger. When considering this option, questions of fraud immediately arise. If, for example, a client wishes to post a buy order at the peak of an asset's price, then (one might think) they could easily stamp the transaction at the most advantageous time, then post it. Or a client might retroactively change the time of an auction bid in order to get in first.

An interesting issue arises, however. The audit trail of a Smart Ledger is strictly chronological by *receipt*. The fact that transactions have a strict chronological ordering is the fundamental property that makes the audit trail immutable. If a transaction is timestamped securely by the client, and there is latency in the network, the order of the transactions on the ledger is not necessarily going to

[23] https://en.wikipedia.org/wiki/Mesh_networking#/

be the same as the order implied by the timestamp. The timestamp is simply another property of the posted item.

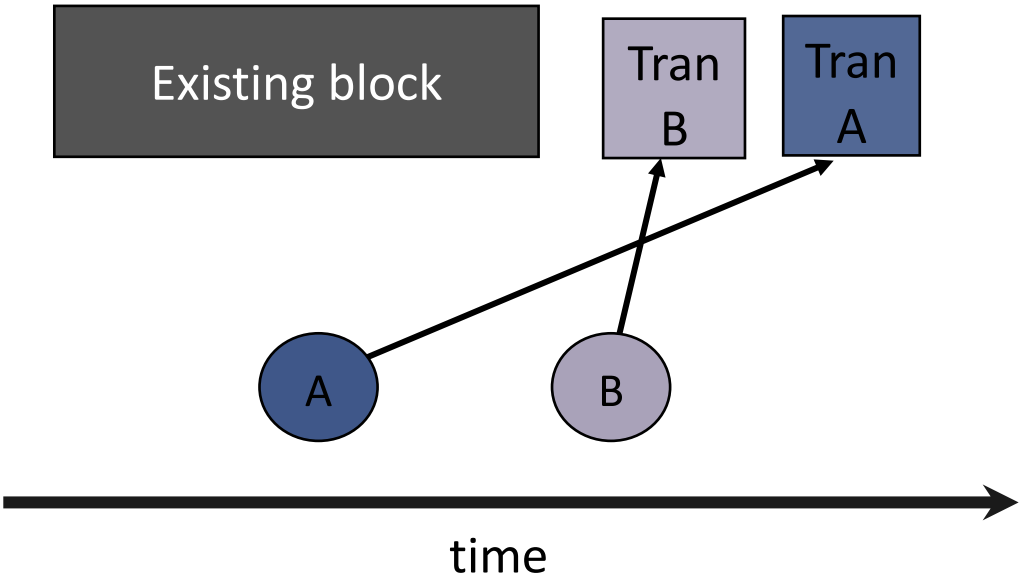

Figure 8: A transaction with a timestamp that differs from the posting time

In the above Figure, transaction A occurs, is timestamped, and submitted to the network before transaction B. However, because the client who submitted transaction A is suffering from high latency, transaction A ends up sitting 'after' transaction B in the audit trail.

Any timestamping system will have a specified precision. It is therefore possible that two events submitted by different users are timestamped so that they appear to have taken place at the same time. In such cases, in order to work out precedence, the Smart Ledger will have to fall back on its chronological audit sequence. The audit trail model ensures that the transactions would have to have been added in one order or another.

B. Time Parameters & Events

Regardless of the mechanism used to stamp the time of *posting*, many smart contracts will at some point be called upon to perform a time *comparison*. At some point a central reference clock will need to be checked against some time variable of the smart contract – a maturity time, a scheduled payment, etc.

The smart contract will therefore have to be able to use time and date values that are set as parameters in the smart contract code, and be able to compare these parameters to a central reference clock. If a maturity parameter, say, matches the current time to a sufficient tolerance, then the dependent event can be fired. Note that it is not necessary for a node running a smart contract to be constantly polling a super-accurate atomic clock signal. When coding the logic, the programmer may simply poll the local clock time, found on the node itself. Most programming languages provide easy access to a machine's local time, through a ***now()*** function or similar. If the network is running a good synchronisation protocol (such as PTP), with a reliable time reference (such as NPLTime®), then the coder may assume that the local clock on the node is regularly corrected as part of normal time synchronisation.

In general, when 'event logic' like this is used in computer programming, the low-level logic that actually waits for the event is often entirely hidden from a higher-level programmer. At the higher level, it is enough to just write a line of code that merely says "wait for *x* to happen", where *x* is, say, "user clicks mouse" or, in this case, "time parameter matches clock time".

But computers run in discrete time, not continuous time. So the concept of 'waiting' in a continuous way does not make sense. The program must instead constantly check and recheck the condition which is being waited for. This is done using a very fast sampling loop.

If the computer is running at more than 1 GHz, that means the period of that loop is of the order of nanoseconds. Every few nanoseconds the maturity time and the current clock time will be compared.

But if both the time parameter and the clock time have a precision of milliseconds, then there is no point in comparing them every nanosecond. If the times to compare are in milliseconds, then it is useless (and wasteful) for the checking loop to run that fast. Typically, the time-checking loop will have a 'sleep' period built in. After every check, the loop will pause for a certain pre-programmed time. Then the loop will restart one millisecond later, and the next check will be performed.

The sleep time is what determines the precision of the time comparison. So if an event needs to be fired within a millisecond of a time being passed, the sleep time needs to be a millisecond (or less). This precision will need to be explicitly specified when coding a smart contract.

When it comes to comparing times, accuracy and precision therefore are intimately related. If times are compared and the comparison is supposedly accurate to the nanosecond, this claim makes no sense if the times to be compared are stored less precisely. In addition, if the two times to be compared are stored to different levels of precision, the more precise time must be rounded before the times can be properly compared.

C. Time Comparability On Smart Ledgers

If the smart contract is running on a Smart Ledger, distributed across a multitude of nodes, how does this affect time comparability?

It is clear that the network running the distributed ledger must be synchronised to the same reference clock. We saw above that using the NPLTime® service, in conjunction with the PTP protocol, is the current industry best practice for time accuracy and synchronisation.

Is such a solution practicable in the case of Smart Ledgers? A key characteristic of Smart Ledgers is that they do not sit under a central authority and have a high level of redundancy. The PTP protocol is designed for a LAN-type setup, with a high degree of control over the network. If Smart Ledgers are to be adopted, they will probably not be geographically or organisationally concentrated.

It is not necessary, however, for every node to be treated as part of the same synchronisation network. Each node could just as easily be part of its own network – say, within an organisation. Provided a node hosting a ledger is using an acceptable synchronisation strategy with an agreed reference, then this will be sufficient. The great advantage of Smart Ledgers is in their use of consensus. If a node is not properly synchronised to the central reference clock, a smart contract running on that node may misfire. But if the majority of the nodes fire at the right time, within a specified tolerance, then the correct behaviour of the smart contract will be executed. Nodes that fire incorrectly will be ignored. There is likely to be a great deal of experimentation on variations of Smart Ledgers using approaches to timestamping and smart contract timing with no consensus, local consensus, or global consensus.

8. Finance & Regulation

A. The Credit Crunch

In 2007, the MiFID directive (now called MiFID I) took effect. Its intention was to create a single market for investment activities across Europe. It was created by European Securities & Markets Authority (ESMA), an EU pan-governmental body based in Paris, for which the UK's responsible body at the time was the Financial Services Authority (FSA). Most of the relevant ESMA functions for the UK are now dealt with by the Financial Conduct Authority (FCA) since 2013.

No sooner had the directive come into effect, than the financial crisis struck. Holes in the MiFID regulation were exposed. It became clear that markets in more complex types of financial instrument (i.e. swaps and other derivatives) would benefit from increased oversight. It was also realised that the question of timestamping had become crucial, as trading had become so fast that it was often impossible to work out from logs the order of trading events. No provision had been made for timestamping in MiFID I.

MiFID I was therefore replaced by MiFID II. Discussing the increased trading in these more complex assets, the EU repeal directive had this to say:

> Previously held assumptions that minimal transparency, oversight and investor protection in relation to this trading is more conducive to market efficiency no longer hold […] [R]apid innovation and growing complexity in financial instruments underline the importance of up-to-date, high levels of investor protection.[24]

In 2012, the UK Government published a foresight report on the question of timestamping,[25] and the recommendations made in the report were eventually incorporated into MiFID II.

[24] http://ec.europa.eu/internal_market/securities/docs/isd/mifid/COM_2011_656_en.pdf

[25] https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/289432/12-1087-future-of-computer-trading-in-financial-markets-summary.pdf

Generally, in financial services regulation the US leads, followed first by the EU and the UK, and then by the Asia-Pacific region. In the case of MiFID II, the EU has taken the lead. In the US, focus on regulation has slipped of late, and the requirements in general are less onerous, particularly in the case of timestamping. At one point the US regulators were proposing a reporting precision of 50 milliseconds – not nearly precise enough to unpick a series of high-frequency trades.

There is a strong impetus for other territories to fall in line with the EU, since MiFID II applies to deals flowing in and out of its jurisdiction. In this way the regulation does end up applying extra-territorially.

B. Markets In Financial Instruments Directive (MiFID) II

MiFID II was published in September 2015 and ratified by the various EU governments. It came into force in January 2018. If the new rules have a theme, it is one of transparency and auditability. These themes can be seen in a few of the goals of the regulation:

- Unbundling research and trading charges,
- Moving markets that were previously conducted over-the-counter (OTC), on to central exchanges,
- Introducing higher levels of surveillance so that market abuse can be detected early and easily, and
- Increasing focus on the buy-side – that is: mutual funds and pension funds who use the banking system to take positions and build portfolios.

The hope is that this legislation will ensure that trading customers are offered the best, most correct prices. Institutions must be able to demonstrate that sufficient research has been performed in order to give the customers the right price. Banks must be able to report trading activity immediately, including price, volume and the timestamp to the specified precision. The transaction reporting requirements have been extended to 65 fields, and institutions are required to store the data for a minimum of 5 years.

While a great deal of attention has been paid to the aspects of MiFID II that mandate the unbundling of research and trading fees, less attention has been paid to the timekeeping and synchronisation improvements that are also required.

C. Article 50 & RTS-25

Article 50 of the MiFID II regulations deals with clock synchronisation. It mandates the drafting of a technical standard to address the problem. This has resulted in the creation of the RTS-25 standard, which addresses the timestamping problem from a few angles.[26]

Firstly, RTS-25 discusses cross-market time synchronisation. It is not enough for the transactions in a single market to be internally consistent. Now, the clocks that govern different markets must be kept in sync. It is hoped that better monitoring across asset classes and jurisdictions will reduce instances of market abuse. If one exchange quotes a price at a certain time, but the clock is running slow, the price quoted will not be the same as the current, correct price prevailing on another, correctly synchronised, exchange. An attacker can exploit this difference by trading an asset between the two markets.

The RTS-25 then deals with precision. An exchange may receive many thousands of orders per second, and the order of their receipt is important. So, it is important that the legislation specifies the minimum time precision for recording reportable events.

It is vital that past events can be reconstructed after the fact. For any given set of transactions a full cross-market report should be possible. Market participants and exchanges must therefore keep all relevant data for all transactions. For this to be possible, each market should not only be using the same reference clock (UTC), but the precision of the time recording should match.

The standard allows for the fact that in some markets, highly accurate recording of events is not relevant or feasible. In the case of voice-based trading, where orders are made and executed on the phone, it is not necessary for the standards of time recording to match those of high-frequency or algorithmic trading. The delays inherent in a human-driven system make overly accurate time measurements meaningless.

[26] http://ec.europa.eu/finance/securities/docs/isd/mifid/rts/160607-rts-25_en.pdf

The fifth point deals with auditability of infrastructure. There is no single specified method that market participants have to follow in order to time-stamp accurately. There are many different ways for an entity to track a reference clock to a specified degree of precision, and synchronise its internal servers. Therefore, the regulators do not lay down a specific infrastructure requirement. Instead, regulatory authorities are given the power to request the details of a market participant's infrastructure. A regulator must be able to understand how market participants are ensuring traceability to UTC, but does not tell them how to do it. Best practice on exactly how to meet these expectations may take some time to evolve.

One major effect of these requirements is that it is now no longer sufficient to have a log of events that are merely internally consistent. Previously, as long as the timestamping of events was done according to some internal clock, it was understood that the times recorded in a market or a jurisdiction might not marry up with each other. This is no longer acceptable. It is now mandatory for most trading venues and participants to agree on an approved reference time, as specified in the regulation.

Many market participants are finding that, for a change, regulation is driving innovation. The encouragement coming from ESMA to define and collect data now means that market participants have a huge source of data to extract analytics from. Much of the data provides valuable insight into the way information flows through an organisation, and could help to streamline trading decisions. For many firms, the changeover to MiFID II in January 2018 has been a catalyst to rethink their systems and even rebuild them from the ground up, rather than continue to adapt multiple internal systems that meet the requirements in different ways.

D. High-Frequency Trading

Even if the credit crunch had not occurred, it was apparent by 2008 that regulation needed strengthening in order to cope with the new challenges of high-frequency trading (HFT). By then, HFT comprised a large fraction of the market. The size of this fraction is not clear. In previous years, numbers as high as 73% have been suggested,[27] but these have been largely discredited.[28] More respectable authorities[29] have put the number at between 10% and 40%, which still allows a great deal of margin of error!

One of the difficulties encountered when estimating the scale of HFT is that of definition. The term 'high-frequency', as applied to trading, is often used interchangeably with the term 'algorithmic'. This is not strictly correct. Equally, the term 'low latency' is used as the inverse of HFT, implying very swift trades.

It is true that high-frequency trading – at the microsecond level – must be algorithmic (or at least computer-driven), since no human can trade that quickly. However, the reverse does not hold. Algorithmic trading is not necessarily done at speed. In general, algorithmic trading can take place at any speed the creator of the algorithm has specified. Which raises the question: exactly how fast do the algorithms need to be, in order to be called 'high-frequency'?

According to ESMA, trading is 'high-frequency' if there are at least two algorithmic trades per second in the same liquid instrument, or if at least four messages per second are sent across all instruments traded over a given European venue. At first glance this seems rather slow to be termed 'high-frequency', but this number would be measured on average over a long period of time – a month, say. There existed proposals to define HFT in terms of median speed across all trading. But apart from being difficult to measure, this would mean that definition would be constantly moving, requiring regulation to move in tandem.

A characteristic of HFT is not simply that orders are placed in quick succession, but that those orders *reach the exchange* as fast as possible. The banks and hedge funds that run the algorithms are trying to take advantage of tiny market imperfections, trying to exploit the inequities before they are arbitraged out. This happens fast, so the algorithms need to take advantage of published prices before they move.

[27] https://www.businesswire.com/news/home/20110608006075/en/Research-Markets-Introduction-Algorithmic-Trading-Basic-Advanced

[28] https://thefinanser.com/2010/07/more-banking-statistics-algo-and-high-frequency-trading.html/

[29] See Aldridge & Krawciw, 2017. Real-Time Risk: What Investors Should Know About Fintech, High-Frequency Trading and Flash Crashes.

Therefore, a system defined as HFT will typically have in place infrastructure that is designed to reduce latency between the system and the exchange. This infrastructure could take the form of:

- Co-location. This is where the institution pays the exchange to host its servers on site, in order to be as close to the exchange's central machines that fill and match the orders. The exchange arranges its customers' machines in tiers. Slots that are physically closer to the exchange servers command higher prices than those further away. The customers' order signals are travelling at or near the speed of light, so the difference in arrival time of each order is negligible. But trading is so competitive, and so fast, that any advantage is felt to be worth paying for.
- Proximity hosting. Another, cheaper, strategy is to rent server space physically near the exchange, in a nearby building. These data centres will either be administered by the exchange itself or by some third party. The clients using the data centres may access the exchange through the same channels as any other market participant, but they hope to gain a speed advantage by being closer on the network.
- High-speed electronic access. Exchanges may offer high-speed access to their servers at a premium rate. This allows customers who are not necessarily located near the exchange to take advantage of market inefficiencies. This means they will hopefully be able to compete with co-located or proximity-hosted customers.

It is worth noting that the kinds of advantages that can be gained through sheer speed have been somewhat eroded. As the sector has become more competitive, it has become less profitable. This has meant that HFT trading volumes have dropped in recent years.[30] More and more market participants are moving to smart trading, using new techniques such as Machine Learning.

Algorithmic trading is not without its dangers. On May 6, 2010, the so-called "Flash Crash" wiped 9% off the Dow Jones Index in 36 minutes. A trillion dollars of market value disappeared, before rebounding equally quickly. It took some time to establish the cause of this, and it is still not clear whether this was human error, criminal malice, algorithmic chaos, or (most likely) some combination of all of the above. What is clear is that during the turmoil, algorithms bought and

[30] https://www.ft.com/content/d81f96ea-d43c-11e7-a303-9060cb1e5f44

sold 27,000 financial contracts with each other, while only actually exchanging a net amount of 200. Clearly, the algorithms were suffering some kind of feedback loop. In the event, many had to be deactivated to try and bring the system under control.

E. Definitions & Differences

MiFID II requires different levels of accuracy for different trading venues. The accuracy required by the Article depends on the gateway-to-gateway latency of the exchange. This is broadly the turnaround time of the exchange in receiving an order and matching it.

If the gateway-to-gateway latency is less than 1 millisecond, the system is considered to be high-frequency venue. For these systems, the precision of the timestamping must be 1 microsecond or less, with a maximum permitted drift of 100 microsecond.

If the gateway-to-gateway latency of the venue is greater than 1 millisecond, the requirements are less stringent. Both the maximum permitted divergence and the minimum permitted granularity are set to 1 millisecond.

The regulation also requires different levels of precision for different types of trading participants. If trading depends on voice-driven or human-based interaction, both the granularity and the maximum divergence are set to one second. For other types of trading – which would include most standard electronic trading – both the maximum divergence and the minimum granularity are set to one millisecond.

Some of the rules laid down by MiFID II are still working their way through the system, and may take a while to fully understand. For example, RTS-25 specifies that certain key events, like “decision to trade” and “order receipt” must be strictly timestamped. But these events are not very rigorously defined. For example, if one wants to timestamp an “order receipt”, one must ask: at what point in an institution’s processes can an order said to have been officially received? Is it at the time the order is displayed on a screen, at the time it is stored or logged in a database somewhere, or at the point it arrives on the fringes of the institution’s network? In a voice trading situation, if the “decision to trade” is considered a loggable event, the question is at what point does the investor make this decision? Since RTS-25 mandates the recording of this moment down to the second, is it the moment the broker answers the phone? The moment the 'Buy' button is clicked? Or some point half way through the conversation?

As the answers to these questions become clearer over time, the extent of an institution's duties will be clarified. When one considers the millions of lines of code which financial institutions run every day, it is highly likely that many of them are in breach in one way or another. The extent of this will become apparent as the process of auditing and inspection begins in earnest.

F. Precision, Accuracy, & Smart Ledgers

At the start of the MiFID II discussions there was a proposal to require timestamping down to nanosecond precision. This was weakened to a precision of the order of microseconds, since anything more precise would place a useless burden on most market participants.[31]

It is easy to become buried in the detail of microsecond accuracy and precision. It is useful to take a step back and think about the goal of the rules, which is to ensure that it is possible to understand, after the fact, what happened and in what order. All the specifications of accuracy and precision serve this goal.

One huge benefit that a Smart Ledger has in this respect is that the entries must be recorded sequentially. Knowledge of the order of events therefore emerges naturally from the technology. The point of the regulation involving 'precision' is entirely to do with ordering, and Smart Ledgers manage it without even necessarily recording time altogether. The precision requirements could almost be said to be made redundant.

The regulation is also focused heavily on reportable events. The regulator must be able to reconstruct a series of transactions. An immutable ledger provides not only a guaranteed order in which the events occurred but, as we have said, it is impossible to alter the record.

[31] Some market participants (including Deutsche Börse) are using CERN's White Rabbit technology, which promises sub-nanosecond accuracy. This is far beyond the reach of most organisations. See http://accelconf.web.cern.ch/accelconf/icalepcs2009/papers/tuc004.pdf.

The accuracy of timestamps, which, as we have said, is required to marry up systems with multiple time sources, is a separate question. There are issues of latency that potentially need addressing. More than that, it needs to be clarified at what point a transaction is timestamped. Is the instant that the transaction is ‘official’ the moment it hits the ledger, or the moment that the client sent it? In most financial systems, the central hub (say, an exchange or a database) determines the moment an event officially arrives. As distributed ledger technology matures, questions of latency will become less of an issue, and then the advantages of the ordered audit will be matched with the required levels of accuracy.

Conclusion

Throughout the history of timekeeping, two important metrics have been vital to the success of any technology: accuracy and precision. Whenever new timekeeping technologies arrive, their accuracy and precision must be evaluated anew.

We have set out the CUTI requirements for a good timestamping system: it must be Comparable, Universal, Traceable, and Immune. Smart Ledgers, properly used, are capable of meeting all of these requirements.

In the world of finance, required standards of accuracy and precision are rigorously set out by regulatory groups such as ESMA and the FCA. This means that if any new technology is to be taken seriously in the financial world it must be capable of meeting these minimum requirements. Indeed, the sequential nature of Smart Ledgers renders some of the precision standards redundant, since the ledgers have to store transactions in the order received. Perhaps the recording of transactions on a Smart Ledger might take some precedence over timing?

However, the distributed nature of the technology does mean that the questions of latency and synchronisation need to be even more carefully considered than a legacy, non-distributed system. This is especially true if the system is to be used in a regulated financial system, such as those supervised by MiFID II.

Principal Author

Sam Carter is a financial services researcher and programmer, with a special interest in quantitative development and natural language processing. He has worked for two decades in the City of London as a developer, quantitative analyst and product manager, in the capital markets and fintech industries. During that time, he has managed and coded the development of systems for reference data, bond and portfolio metrics, synthetic collateralised debt obligations (CDOs) risk calculations, commodities pricing, and automated compliance. He holds an MSc in Financial Mathematics from King's College London, where he worked on credit default swaps, and an MA in Linguistics from UCL, where he worked on information theory in the syntax of natural language.

He is interested in the interfaces between the different actors in the financial world – the dialects spoken by people in finance, mathematics, legal, compliance and trading – and what information is lost when those worlds attempt to communicate with each other. Sam is an amateur guitarist and pianist. He speaks terrible French and even worse Afrikaans.

Acknowledgements

The author would like to thank Makoto Seta of the FCA, Leon Lobo of the National Physical Laboratory, Ian Salmon of IgniteG2M, and Michael Mainelli of Z/Yen Group for their patient explanations and contributions.

Distributed Futures is a significant part of the Long Finance research programme managed by Z/Yen Group. The programme includes a wide variety of activities ranging from developing new technologies, proofs-of-concept demonstrators and pilots, through research papers and commissioned reports, events, seminars, lectures and online fora.

Distributed Futures topics include the social, technical, economic, and political implications of Smart Ledgers, such as identity, trade, artificial intelligence, cryptography, digital money, provenance, FinTech, RegTech, and the internet-of-things.
www.distributedfutures.net

Cardano Foundation is a blockchain and cryptocurrency organisation based in Zug, Switzerland. The Foundation is dedicated to act as an objective, supervisory and educational body for the Cardano Protocol and its associated ecosystem and serve the Cardano community by creating an environment where advocates can aggregate and collaborate.

The Foundation aims to influence and progress the emerging commercial and legislative landscape for blockchain technology and cryptocurrencies. Its strategy is to pro-actively approach government and regulatory bodies and to form strategic partnerships with businesses, enterprises and other open-source projects. The Foundation's mission is the promotion of developments of new technologies and applications, especially in the field of new open and decentralised software architectures.
www.cardanofoundation.org

"When would we know our financial system is working?" is the question underlying Long Finance's goal to improve society's understanding and use of finance over the long term. Long Finance aims to:

- expand frontiers - developing methodologies to solve financial system problems;
- change systems - provide evidence-based examples of how financing methods work and don't work;
- deliver services - including conferences and training using collaborative tools;
- build communities - through meetings, networking and events.

www.longfinance.net

Z/Yen is the City of London's leading commercial think-tank, founded to promote societal advance through better finance and technology. Z/Yen 'asks, solves, and acts' on strategy, finance, systems, marketing and intelligence projects in a wide variety of fields. Z/Yen manages the Long Finance initiative.

Z/Yen Group Limited
41 Lothbury, London EC2R 7HG, United Kingdom
+44 (20) 7562-9562 (telephone)
hub@zyen.com (email)
www.zyen.com