

Dated : 14 January 2015

To,

**Mr. B. K. Gupta, DGM**

Market Regulation Department

Division of Market Supervision

SEBI Bhavan,

Plot No.C4-A,'G' Block, Bandra Kurla

Complex, Bandra(East), Mumbai -400051

**Sub : Market manipulation using technology at National stock exchange(NSE) collocation datacenter across 3 years(2011 -2014) in Derivatives with active collaboration of NSE data centre staff. Information from a whistle blower.**

I wish to draw your attention to a sophisticated market manipulation done at NSE for several years at NSE collocation. The most shocking aspect is that when the matter came to knowledge of NSEs management team they have chosen to hush up the matter under the carpet rather than coming out in open against the same.

Since this has been a highly complex front running system it will require some level of technical expertise to understand how the scam has continued so long with sebi (and even NSE for quite some time ) not realizing what was happening. Before I explain the modus operandi to you I would like to inform you of my background so that you would be able to better appreciate the source of my knowledge. I work in the technology team for a Singapore based hedge fund and we have a significant exposure to Indian stock markets and therefore we analyze our trade quality etc.I have tried to keep all explanations as simple as possible , but still if you are a non technical person I would suggest it would be a good idea to read this letter with a IT geek around to help you map the system.

The market manipulation I am referring to has been occurring by enabling certain vested brokers to get market price information ahead of the rest of the market and thus enabling them to front run the rest of the market. To understand how this works you will first need to understand how price information is disseminated by NSE. NSE has two price feed streams one is the broadcast (UDP stream) which is used by NSE for its VSAT network and leased line networks .This multicast (multiple broadcast which everyone view concurrently like a TV) is limited to 2 mbps due to NSE's leased line network capacity. The other price feed is called TBT (tick by tick)<sup>1</sup> which is actually reflecting every new order price information coming into the exchange system. Due to the large size of the price information this is available only at NSE collocation. The challenge to NSEs tech team was they felt they could not

generate this stream as a broadcast (though now they have upped their technology and have started this a few months back and is called Multicast TBT in the present avatar). As a result the TBT information was disseminated over TCP/IP (in this mode its like information is delivered one by one instead of a broadcast). This is the crucial point which needs to be appreciated fully. In a broadcast environment everyone gets the price information at the same time (assuming they are at the same distance from the server- which is a given at collocation). The life in TCP/IP is very different though. In TCP/IP since the processing of price information is sequential as everyone has to be given this information and it cannot go concurrently, the person who gets the price information first would be able to process it first ahead of everyone.

When NSE started TBT in 2010 all ISVs wrote programs using NSEs API (application program interface- which defines the protocol as per which information will flow from the exchange). Logically all people should be on equal footing. The issues started since this was essentially a patch work solution put together by NSE it had several internal inconsistencies. The system had to cater to a significant load which it was not architected for. One TBT server of NSE had to service several connects ( i.e. servers of the different members). Since the price information available sequentially it simply meant that the person who connected first to the server would get the price information ahead of the others every time throughout the day. E.g. If I was first to connect to the TBT server in the morning and subsequently 40 other guys connected to the server and the exchange sent out a price packet it would first send to me and then in order of connection sent to the other 40 guys. This cycle would be repeated for each and every price information sent by the exchange make me ahead of the other 40 guys throughout the day. In theory there is nothing which can be done to bring everyone at par in TCP setups that's why globally UDP is the preferred way of price dissemination. The only redeeming assumption is next day I would not be the first guy so across days the advantage would even out across time. This is a crucial point to appreciate as this forms the crux of market manipulation.

Further complications in the TBT stream arose as demand for collocation rocketed so the exchange had to put up multiple servers to all allow different client servers to connect and listen to the feed. This meant that the load was not necessarily same across the servers. So if I could know the server with lesser load I could connect to the same and hopefully be ahead of the crowd. If server 1 had 20 clients and server 2 has 50 clients, naturally all those connected to server 1 would be faster than those connected to server 2. By the nature of technology I could not hop across servers intraday as when I go to connect to a new server I would always be the last guy on the new server and hence slower than everyone on that server. **It thus became crucial to be first one connect to a lowest load server.**

Life becomes somewhat more technical from here. Exchange needed two more things to give price information consistently. One it needed to ensure there was no skew in the load across servers (ie everyone is not one server and second server is idle). Secondly it needed some form of backup arrangement in case its server failed so it introduced a back server. Everyone connected on the

primary server would be switched over to the backup if the primary failed seamlessly and the backup server was always running for business continuity. The implications of both these aspects will become clear to you as you understand how the manipulation was done.

NSE came up with a simple system to handle the load skew. It simply mandated the server to which each of the clients servers could connect and correspondingly the back up server. By this it would know upfront of the number of connections on a server and try to keep the load constant across servers apart from planning capacity to handle the load and of course invoice its members. If anyone tried to connect to a different server NSE would be able to track and would call up the member and ask him to fall in line.

I must take a pause and brief you about the people who were in charge of the design and execution of the above system. During the period the data centre and collocation came up Ravi Apte was the CTO of NSE and the person overall responsible for the design of the system. On the data centre front Jagsish Joshi was the person in charge. In all senses he was the person who controlled all the ground level nuts and bolts at the NSE data centre. At this time it is also essential to understand the ISV space . Omnesys was the market leader with its DMA product being highly popular on the institutional desk. Its algo trading solution was also popular among the domestic prop firms. It is important to note here that since NSE was the second largest shareholder of omnesys and worked closely with them for several projects like NOW , omnesys and NSE staff was on better terms than NSE was with any of the other vendors. The other major domestic ISV was greek soft. Among the global players sunguard was the leader in DMA solutions with flextrade and RTS being the major players on the Algo front. Rest of the solutions were either built out of components like RMS and OMS taken from one of the above vendors or built entirely in house.

While the collocation was started by NSE in Jan 2010 it took some time for the more enterprising to figure out what was happening where in the system . By the end of the year the smart guys had all figured out that the way to game the system lay in being the first one to connect to the server and preferably a server which was the fastest . A server could be faster due to lesser load (20 v/s 50 guys listening) or it could be that the hardware of the server was slightly more powerful.

At this stage S G of securities a NSE member entered the fray . He was a non IT person very good at building informal relationships with people. During his interactions with people in the data centre team it became clear to him where the advantage lay. He needed a person who knew the nuts and bolts of the system. He was working on omnesys for his algos so he took the person in charge of the nse collocation installations, a person called nag on his payroll. With this team he would attempt to figure which server was performing better and gradually he warmed up a few persons in the data centre team to let him know the time at which the servers would be started so that he could



be the first to connect. For some strange operational reason (the motives thereof will become more apparent as you read further) NSE did not start its TBT servers at a fixed time every day. Firstly it could not physically start all TBT servers concurrently and unlike the market which opens at the same time for all persons the TBT connection was established before the market open (typically one hour before) but the time was not fixed. This was the first of the many steps which cemented the bond between data centre staff and OPG. Every day he would be privy to the information as to which server would be started at what time so that he could be the first one to connect and enjoy the advantage.

Over a period of time others figured out the importance of this and did what all geeks do. They wrote a piece of code which would continuously ping (ie check if the TBT server was started) and attempt to connect immediately when it got the first response from the TBT server. In fact Omnesys launched a managed data centre service for its algo clients where it would technically be managing the infra (Servers) as well as software for its clients. It claimed it provided hardware and wrote custom software for the same under managed services and as a result had far superior performance to any other algo system around. The crux of this was the knowledge that connecting faster would put the server ahead in the queue. In fact it became so popular that Omnesys started 20 % profit share from its clients for the profits made on the system (why it resulted in profits always will become clearer further down the letter). This was obstinately billed as manhours for managed services so that it could be variable every month. This was a very clean way for them to earn out of the knowledge. Everyone just believed it was a difference in their software which gave them an edge where as it was simply front running.

Once omnesys started doing this on a large scale S G felt he was losing his competitive advantage. He needed to figure out some way to keep his edge. He also realized using a mass market product like omnesys the facts would eventually leak out somewhere so he went ahead and built a in house software which no one else would have access to and therefore no one would be any wiser of what he was doing. He also realized efficiency of the software was not of prime importance due to the kind of advantage he enjoyed in connection. To understand the magnitude of the opportunity you need to put the timeframes on a relative scale to see the kind of economic advantage reading the TBT feed leads to.

The order of magnitude it takes to process a price and throw out an order to the exchange typically ranges between 50 to 200 micro seconds (difference between an efficient and inefficient system). This typically means one fifth of a millisecond (1000 micro seconds is one milli second – 1000 milli seconds makes one second). The time taken at nse to receive and order and send out confirmation ranged from 2 to 10 ms (milli second). This was variable and is directly proportional to the load on the exchanges systems. The average time has been coming down from 10 ms in 2010 to 2 ms in 2014 as NSE has improved its technology. The order of variation between different TBT servers ranged from 5 ms to 50 ms, which was a function of the load on the TBT server. Further variability would come on basis of the sequence number which you connected to a given TBT. The order of magnitude would now

make it apparent why order processing speed was not important as the ability to game the system. If you could connect first to the lowest load server every day you would be 100 times faster than the average person on collocation. About non collocation persons well they were the people who were never even reached the stadium where the race would take place.

In terms of the economic advantage in being ahead by 20-50 ms (which is just one fortieth of a second) you would need to understand what HFT firms do globally to make riskless profits. They simply hedge the price inefficiency in a way that when the price normalizes they would be able to exit very easily with a profit. At NSE the only way to achieve maximum profits using this short time advantage turned out to be an order type called Multilegged IOC (immediate or cancel). A multilegged IOC order is an order with 2 or 3 legs (sub orders) which get filled (i.e. traded) if all the orders are getting executed simultaneously. If price for one of the orders is not available all the orders get cancelled by the exchanges system. A simple example would enable you to see how one can profit from this. Suppose stock X has current month future quoting at 100 /100.50 and next month is quoting 101/101.5. Now suppose for some reason there is a large sell order which comes and fills all buyers till 99.5 and becomes best seller @99.5 so that the book may look like 99.25/99.5 (10 lots). If you had this information 20 ms ahead of the crowd all you need to do is to fire a 2 legged IOC order to buy 10 lots @99.50 and simultaneously sell 10 lots in next month till 100.75 (ensuring an assured profit of atleast 25 points). The beauty of the system is that if there are buyers for 5 lots in the far month till 100.75 the system will match order for 5 lots and cancel rest of the order. Once the market reverts to its mean you can close out the open position and make 25 bps riskless profit. On NSE the traders apart from doing the above 2 legged trade also did the 3 legged conversion reversion strategy which is a way of capturing riskless profits of mispricing in NSE options. In this when the future is purchased a call is sold and a put is bought (of the same strike price). I trust you can figure out the mechanics easily after the above example or you probably are already aware of this kind of a trade. In the above example you would have also understood why this is a winner take all market. Once the first IOC order has got filled all the orders behind in the queue will simply not get anything. This may also now explain to you the basic reason why SEBI had to impose an order to trade ratio penalty system to dis-incentivize order flow in the derivatives segment. Every time there is a mispricing all the servers fire IOC orders in a hope to capture the mispricing and only one gets the fill.

The quantum of economic advantage is actually linked to the reason why these trades occur and relative to the rest of the market where the system stood in terms of knowing the price information. The first kind of advantage typically occurs when someone is trying to exit or enter into a large position quickly. as a result the buyer or seller tries to hit the current order book and once the existing orders in the book get filled the order becomes a resting order leading to a mispricing. The second more windfall gain occurs when there is a fat finger or a punching error. In such cases say if someone enters the incorrect price in an option strike the person reaching the system first wins riskless money. You may be wondering how much would the system be worth. Well taking the example of securities

alone I could tell you that it was worth approx 2 crores per month atleast or roughly a 100 crores in three years (though the balance sheet of Indian firms never seem to indicate such figures due to tax juggleries). The quantum in case of Omnesys is easier to guess you simply need to know the figure of managed services invoicing and multiply by 5.

Now returning to the sequence of events on the ground as they transpired. OPG securities figured out that more people like omnesys had figured out the importance of connecting faster and they therefore needed some way to remain ahead of the curve. This is where NSE collocation staff proved specially helpful. They would grant him access to the time they would switch on the servers and would give him access to servers which were latest and least crowded. The consideration thereto would be obviously off the record. Officially he would always claim the superiority of the software but the backend was always this. The other market players were not able to ever decipher due to one more important piece which was surprisingly but crucially missing from the NSE collocation. There is globally a practice of time sync across all exchange servers and all servers at collocation and sync themselves with the exchange clock (down to milli second). Though the device costs only a few thousand dollars it was surprisingly missing from the exchanges infrastructure. In effect it meant no one could prove latency numbers as every computer has its own timestamp. This advantage remained a solid lead for OPG all thorough out 2011 and later through most of 2012

Around September 2012 there was a change of guard at NSE . Ravi apte retired as the CTO. In his last days his level of interest in the affairs was very limited and much of the control had shifted to the operational staff at the data centre. As a passing bird he needed to pay lip service to complaints of brokers of problems at collocation. When Umesh Jain joined as CTO in oct 2012 he started off by meeting a lot of players in the system and when he discovered this most common issue was TBT servers he setup a simple solution in a couple of months. Being a new person willing to listen he found himself flooded with logs indicating the high differential between multiple servers. Latency difference of 50 ms across servers was common. He realized this needed to be fixed. He dismantled the system of TBT allocation and to ensure fairness put in what is called a load balancer. The load balancer would be the server to which everyone would need to connect and in turn it would dynamically decide which TBT server the application would connect internally based upon the overall load so that load is evenly distributed.

This change clearly represented the biggest challenge to the model of those who were trying to game the system. Omnesys suddenly found it had lost the edge in IOC servers and saw a drift away from its managed services. OPG then figured out there should be a way to outsmart the system. Once more the obliging data centre staff came in handy. While the system had an automated allocation of TBT servers they still had the backup servers which were to be connected only in case the primary went down though they were equally powerful. They figured out if they let OPG connect to the Backup servers with zero load he would still be getting far better latency. They obliged and the cycle continued.



Around this time coincidentally a couple of global players went live with their servers at NSE collocation. The local guys could not figure out the working but the global players clearly had more exposure of what was happening. One more improvement which the new CTO brought was the ability to sync servers with the exchange clock. Suddenly the cat was out of the bag (for a select few who understood what happening). It was now possible to see what time the exchange generated the price broadcast, what time it was received by the different servers and what time the trade got executed. Since it often got filled before the broadcast was received by the market it was obvious something was amiss. Exchange simply could not explain it away as random differences or the abilities of the different software's. Armed with the logs (though not necessarily understanding why this was occurring) different firms started approaching NSE to explain the sequence of events. While officially the stand was variability of TBT servers and loads, someone within NSE must have started understanding what was happening. NSE decided to go to the global standard of multicast at collocation so that everyone would get the price concurrently.

In hindsight the events which unfolded clearly leave a trail of what most probably happened. Someone smart in NSE must have put the complaints and the trades together. Suddenly there was a shuffling of lower level staff at the NSE datacenter. Practically all the operational people were changed. The issue of handling senior staff was still there. If they were fired they could come out in the open and if they were not moved they would again create some similar situation. The matter was quietly hushed up. And everyone was given an honorable exit. The external parties who understood the happenings were quietly accommodated in different ways (I will come that near the end of my letter) and in a surprising move the data centre incharge left to join MCX SX as data centre incharge (after the Jignesh shah episode) so its not difficult to add up things. NSE soon started multicast at its colo. Suddenly market share of OPG fell off the charts and omnesys managed services went off the track. (by now NSE had also sold off its stake to reuters so it had no specific interest in the future of the vendor).

I know your concern may be what if this is just a cock and bull complaint by some vested interest. Well frankly I have some vested interest which I will come to later, but the proof of the pudding lies in eating it. The easiest way to verify or throw out this complaint is to go back to the data and see if it supports the same. If NSE stores its internal order logs of order and trade with milli second timestamp, it will be very easy to verify when did the first trade happen which caused an abrupt price movement and when did the IOC multileg order get fired by different members. If the latency difference between the orders generated across members is in the range of 200 micro seconds there is obviously no foul play, but if the difference is in the range of 20 milliseconds you know the working explained above. As a confirmation you may verify that this pattern exists across days or is a one off event for the member. In case the millisecond time stamp is not available for some hard to fathom reason NSE maintains an order number which is chronological in order of the receipt of order by its system. For the purpose of analysis you need to filter all the multilegged IOC orders received by the exchange on few sample volatile days in say 2012. If the order numbers of a selected few members are always ahead of similar

orders received from other members you can be sure of the front running. You can also examine the fill ratio of orders across members (ie how many of the IOC orders fired got traded as a percentage). Since the working has been explained in detail above you should be able to correlate the same with the above data. One more easy way to cross verify would be to check which member server connected to which TBT ip of the exchange as it would clearly indicate the load skew, but this may not be available if NSE is not regulatory bound to store this information. One alternative may be to check with global vendors if they have stored any such logs .

- Coming to my part of the interest. I was actually working for one of the Singapore based firms and have shifted to a new HFT firm recently. The problem which we face is that in quid pro quo for keeping quiet NSE has given certain privileges to my earlier firm which are not available to me at the new firm. This is resulting in degraded performance of the algos. If SEBI will take some action I am sure the current unfair advantages would stop as NSE will be forced to be more cautious otherwise it will be difficult for us to work in an unlevel playing field. By the advantages I am not referring to procedural benefits like allowing customized racks or giving bigger order message lines , I am referring to structural advantages which are more difficult to prove with data else I would have pointed the same out as well (after all NSE has also learnt as part of the previous fiasco) .

I need to protect my identity currently as I am unsure of the kind of protection I would get as a whistleblower. Technically since my previous contract was based in Singapore I may not be able to protect myself from the contractual clauses till a malpractice is confirmed from your end and hence I am apprehensive in disclosing my full identity currently. If at a later date due to prosecution requirements I need to come forward openly I shall be willing to do the same. Currently I am unsure whether SEBI would be taking this matter forward and prosecuting the guilty or this will die as a faceless complaint. I will be updated once SEBI initiates any significant investigation. I will be able to provide the logs indicating the differences in performance across TBT servers at NSE colo if that is of any help in the investigation. In order to identify myself at a later stage I will confirm my name as \_\_\_\_\_ and the address at singapore as \_\_\_\_\_

Yours faithfully

✓ C.C.- **SUCHETA DALAL, Moneylife Foundation**, 304, 3rd Floor, Hind Service Industries Premises,  
Off Veer Savarkar Marg, Shivaji Park, Dadar (W), Mumbai – 400 028

\_\_\_\_\_ **CONFIDENTIAL INFORMATION TRUNCATED**