OnePipe 3.0: The Next Generation Dark Liquidity Aggregator

It is currently estimated that 15% to 20% of the volume in the US cash equity markets is executed in dark pools or using hidden orders.[1] As such, interacting with dark liquidity has become an essential part of the buy-side trader's workflow. In late 2007, Pragma pioneered the idea of a dark pool aggregator with Onepipe, a trading algorithm that connects to and searches for liquidity in multiple dark pools simultaneously. Using information from its own orders and executions, OnePipe adapts its order submission strategy to capture as much liquidity as possible. As dark pools have become a larger segment of the market, and a wider variety of market participants now interact in the dark, the need for more sophisticated dark pool aggregators has increased as well. In this paper we introduce the latest version of our award-winning dark pool aggregator, OnePipe 3.0.

……

I. Summary

OnePipe 3.0 is the latest version of Pragma's dark pool aggregator, and includes significant improvements in allocation logic, dynamic use of minimum fill sizes, and support for customized trading goals. The new allocation logic, through more strategic sizing and placement of child orders, can improve cross rates by 20% relative to the previous allocation scheme. In addition, by dynamically adjusting the minimum fill quantity of child orders, OnePipe avoids small fills when the amount of liquidity available in larger blocks is sufficient to satisfy the trader's goal, but permits interaction with small orders when necessary. Finally, the new OnePipe allows the trader to express his trading goals explicitly as a participation rate.

II. Dark Pool Aggregation - Historical Survey

As the prevalence of dark liquidity and the fragmentation of the US equity markets have increased, dark pool aggregation has emerged as an essential algorithmic trading tool. A dark pool aggregator is a trading algorithm that connects to multiple dark pools and distributes pieces of the parent order among these pools. As child orders are executed in the various pools, the algorithm re-allocates the remainder until the order is fully executed.

The first generation of dark pool aggregators used a ``bathtub'' approach. The bathtub approach divides the parent order equally between the various pools. As executions occur, the algorithm shifts the quantities from one pool to another in an attempt to keep the same quantity at each pool. While this approach can provide reasonable performance, it is not optimal under realistic conditions. In particular, the bathtub approach assumes, implicitly, that all the dark pools are the same, and that the flow is uncorrelated. However, as we know, these two assumptions are false. For example, the type and volume of liquidity available in SigmaX, Goldman Sachs' internalization pool, and BIDS, a consortium-owned block pool, are quite different. In addition, as our research demonstrates, there is a significant correlation in the destination in which consecutive executions occur.

Examining a simple example demonstrates the drawbacks of the bathtub approach. Suppose the aggregator is connected to 30 destinations, and there are 3,000 shares left in the order. A dark pool aggregator that is based on the bathtub approach will submit thirty 100 share child orders to each of the pools. However, if one of the pools is much more active, say Nasdaq Mid-Point, then we are almost guaranteed to have a lower trading rate, as much more liquidity is likely to pass through Nasdaq Mid-Point than a small dark pool. A better allocation would place more volume at Nasdaq Mid-Point in order to reflect the higher probability of liquidity arriving there.

As market participants realized the disadvantages of the bathtub approach, the second generation of dark pool aggregators began to emerge. The second generation tried to improve on the main disadvantage of the first generation, namely the lack of feedback between where executions occur and the allocation of child orders to the destinations at which most of the executions occur. The second-generation aggregators used ad-hoc mechanisms for adjusting the order size in response to a fill in a particular pool. In particular, these algorithms would typically overweight a "hot" pool as executions occurred at that pool, and would revert back to equal weighting as time passed with no executions occurring.

While the second-generation aggregators resolved many of the problems of the first-generation dark pool aggregators, new problems emerged. In particular, as dark pool aggregators became more efficient in reacting to liquidity and adjusting the allocation to reflect the true distribution of liquidity in the market, their trading rates increased. As a result of the increase in trading rate, and as the range of participants using dark pools grew, aggregators became more exposed to issues like gaming and information leakage.[2] The increase in trading rate could be used by other market participants as a signal to front run the order. In response, various techniques for protecting the order emerged, such as setting static minimum fill sizes to reduce the trading rate and protect the order. These solutions were mainly ad-hoc and affected the orders in some unanticipated ways. For example, at times, cross rates might have been very low even in the presence of significant "safe" liquidity in the market trading at sizes below the minimum fill size threshold.

---

[1] "Let There Be Light," Rosenblatt Securities's Trading Talk.

[2] "Anti-Gaming in the OnePipe Optimal Liquidity Network," Pragma Securities.

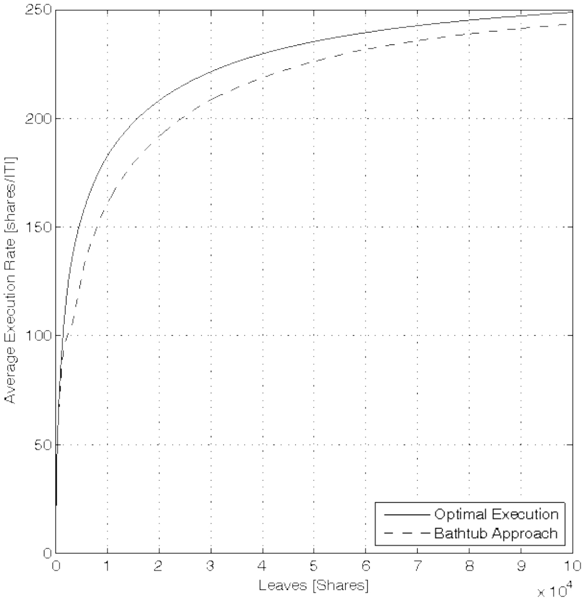

Figure 1: Expected cross rates per for the bathtub and optimal allocation approach.

III. Third-Generation Dark Pool Aggregator

OnePipe 3.0, our third-generation dark pool aggregator, provides a coherent solution to many of the problems identified in the first and second generation of dark pool aggregators. This new algorithm is designed to source liquidity more intelligently, while adhering to explicit trading goals set by the user.

III.A Optimal Allocation in Dark Pools – Block Trades

The optimality of a dark pool allocation methodology can be measured by the amount of liquidity that is sourced - the more the better. Under this criterion, an optimal allocation of an order in dark pools will maximize the expected trading rate or equivalently, minimize the time to complete the order. As previously mentioned, for large orders, the bathtub approach is almost optimal, but as orders get smaller this approach breaks down, and better allocation methodologies exist. The optimal solution, one that maximizes the trading rate and can be shown to depend on the distribution of the number of shares to arrive at each pool.

Figure 1 depicts the difference in expected crossing rates of the bathtub and the optimal allocation as a function of the order size. As can be seen in the figure, for very large orders, the crossing rates of the bathtub approach and the optimal allocation are almost identical. However, as the number of shares decrease, the difference between the bathtub approach and the optimal allocation scheme become apparent. Improvements of up to 20% can be achieved by using the optimal allocator.

Unfortunately, the equity markets are not only fragmented but also very dynamic. Liquidity begets liquidity and as a result liquidity flows from one market center to another. OnePipe continuously adjusts its estimate of the distribution of liquidity to arrive at each pool and allocates between the pools based on this evolving distribution.

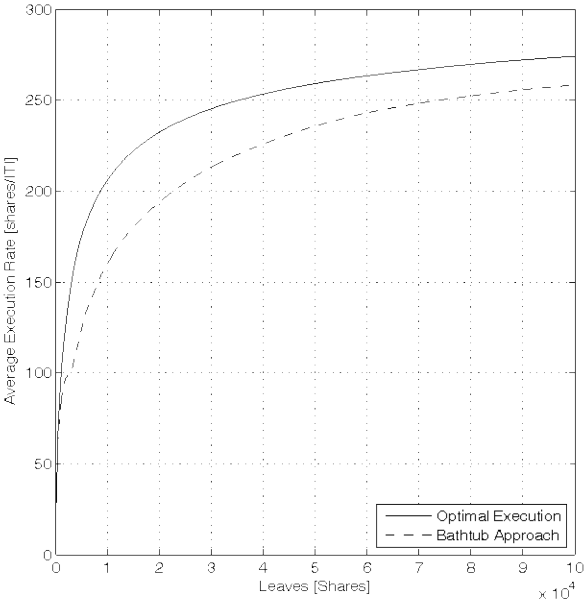

Figure 2: Expected cross rates for the bathtub and optimal allocation approach. A destination that is responsible for 1% of the volume on average was adjusted to account for 30% of the volume.

Adjusting for this phenomenon is critical in maximizing the trading rate. In Figure 2, we depict the expected crossing rates for both the bathtub and the optimal allocation as a function of the order size. However, in this figure, in contrast to Figure 1 we artificially assumed that a destination which is normally responsible for 1% of the volume is now responsible for 30% of the volume. Figure 2 demonstrates the importance of using the optimal allocation over the bathtub approach.

Additionally, much intelligent behavior falls out of this solution of estimating the full distribution of liquidity at each venue. For example, block oriented destinations like PULSE or AQUA are accessed by periodically sending an order for a short period of time to the pool, probing for latent liquidity. The period is based on the historical fill rate in the pool, recent fills if any, and market conditions.

III.B Minimum Quantity and the Dark POV Concept

An evolving estimate of the liquidity in each pool allows OnePipe to achieve explicit trading objectives by setting the minimum execution quantity for each child order. For example, assume that our estimate points to an expected trading rate of 127 shares per second. If the current order has a large leaves size, and we are using optimal allocation, we will achieve, on average, a trading rate of almost 127 shares per second. However, if we set the minimum quantity of each child order to say, 300 shares, we will miss every order for less than 300 shares, and we can estimate our trading rate will be lower than 127 shares, say 87 shares per second.

The concept of dynamically setting the minimum quantity to achieve a target trading rate is very powerful and is at the core our new implementation of OnePipe. For example, suppose the trader would like to achieve a trading rate of 20% in dark pools, but still have the option to finish the order if a "natural" arrives at a certain pool. OnePipe will set the minimum fill quantity to achieve a 20% participation rate with the largest fills possible. Suppose, based on our calculation the minimum fill quantity is set to 300 shares. If a series of orders that are 1,000 shares or more arrive into the pool, OnePipe will shift the remaining quantity into this pool, quickly completing the order. On the other hand, if there is a significant amount of liquidity in the form of 100 share orders and no orders larger than 300 shares, OnePipe will reduce the minimum fill quantity, allowing it to participate with the available liquidity and blend with the flow. Thus, OnePipe adapts dynamically to the trading characteristics of individual stocks to ensure a more uniform execution experience for traders.

III.C Anti Gaming and Lifeguard

It is by now well understood that trading in dark pools has the potential to impact stock price just as much as trading in the lit markets, and one of the mechanisms of this impact is gaming and front running by predatory market participants. The liquidity in dark pools can be divided into natural and opportunistic liquidity. Natural liquidity consists of orders submitted by directional traders who generally plan to hold the position for longer than one day. One important class of opportunistic liquidly is "pings" - orders generated for the sole purpose of detecting the existence of a larger institutional sized order which might then be gamed or front-run. Often opportunistic liquidity is referred to as "toxic" and is associated with large impact. While natural liquidity is the vast majority, a small portion of dark liquidity is toxic, and countermeasures to protect execution quality are an essential element of any high-quality dark pool aggregation algorithm.

OnePipe's dynamic allocation methodology is a significant improvement in our Lifeguard anti-gaming logic to protect orders against toxic liquidity and combat information leakage. Toxic liquidity usually consists of smaller orders, 100 to 300 shares, as speculators must balance risk against the value of uncertain gaming opportunities. For this reason, setting a static minimum fill quantity has long been a tool to protect orders against gaming. However, a statically set minimum fill quantity may be too high, resulting in low crossing rates in situations where gaming is not an issue. In other situations, a statically chosen minimum fill size may be too low, resulting in significant impact. OnePipe 3.0 dynamically sets the minimum fill size at the highest level that will support the desired trading goal. When a significant amount of liquidity arrives in the form of large orders, the algorithm will increase its minimum quantity, which in turn will reduce the footprint the algorithm has in the market, and when liquidity is available only in small orders, OnePipe will adjust accordingly.

The use of dynamic minimum fill quantities for the child order reduces the probability of being detected by a gamer, but a speculator might still infer the existence of an investor's large order, for example by examining the consolidated tape in real time. As another layer of protection, we use dynamically adjusted limit prices on each child order. Lifeguard, applies two layers of protection in order to determine the optimal limit price for the child orders. The first layer uses short-term price moves in order to determine a fair price and a band around it. If the current price is unfavorably outside this band, the algorithm will not trade. In particular, this mechanism protects against gamers who push the price faster than expected natural price movement. In addition, by anchoring the limit price to a point in time in which we considered the price "fair," we achieve another layer of protection. Together, we use the most conservative limit price of the two to set the limit price for the child order. This dynamic limit price continuously evolves based on current market conditions, and allows the algorithm to avoid being executed at unfavorable prices even if the order is detected by a gamer.

IX. OnePipe 3.0

OnePipe 3.0 enjoys all the benefits of the optimal allocation methodology discussed above. Five urgency levels of OnePipe target different trading goals, each expressed as a desired minimum trading rate to be achieved in dark pools. The algorithm dynamically estimates the amount and type of liquidity in each pool, and based on these estimates builds an allocation that uses the highest minimum fill quantity possible that is likely to achieve the desired trading rate. The algorithm has a feedback mechanism used not only to increase the minimum quantity but also to decrease the minimum quantity when necessary. Note that the algorithm will increase its minimum quantity not as a response for trading too fast or too much, but as a response to a change in the distribution of available liquidity, allowing the algorithm to capture natural liquidity while avoiding toxic small orders.

In Figure 3 we depict the general structure of the algorithm. Fills are fed into a liquidity estimator that estimates the amount and type of liquidity at each pool. This algorithm combines historical information and the real time fills to adjust for the estimated liquidity in each pool and the market as a whole. This information is fed into the allocator, which uses the estimated liquidity and fill information to determine the allocation that has the highest minimum fill quantity likely to achieve the trading objective. Finally, the allocation is fed into our anti-gaming logic to generate child orders with the appropriate limit prices, and sends them to the various pools.

Overall OnePipe 3.0 solves many of the problems discussed in the previous sections. Its optimal allocation algorithm maximizes the crossing rate given the leaves of the order. The algorithm dynamically changes the minimum fill quantity of child orders to avoid toxic liquidity when possible, while ensuring continuous interaction with the market. Finally, the addition of dynamic minimum fill quantities represents a significant improvement to our Lifeguard order protection logic.

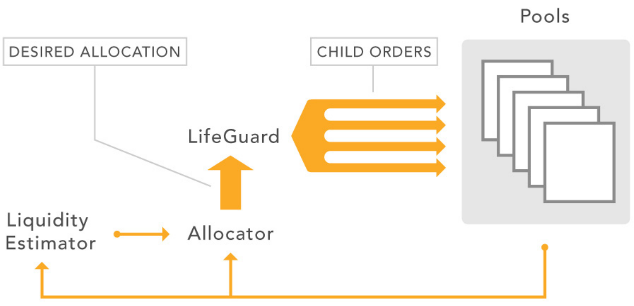

Figure 3: OnePipe 3.0 block diagram

For questions or comments, please email Eran Fishler, Director of Research (efishler@pragmatrading.com), call 866.502.2546 or email algosupport@pragmatrading.com.