traderserve

“Breaking the Barriers”

A Technological Study of the Obstacles to Pan-European Best Execution in Equities

Released March 2003

Nicholas Hallam and Nick Idelson
Traderserve Limited
12-16 Laystall Street
London EC1R 4PF

Preface

This report is based on interviews and private communications conducted with execution venue connectivity providers (ISVs), front office system vendors, exchanges, data consolidators and vendors, transaction cost analysis experts, clearing and settlement houses, retail and institutional equity brokers, retail service providers, standards organizations, fund managers; the literature cited in the bibliography; and on the authors' own experience in the trading, fund management and technology provision arenas.

The equities world is subject to very fast change at the moment. This affects execution venues and clearing and settlement providers alike. In particular, there is very rapid change in the US involving alternative execution venues, and in Europe concerning Central Securities Depositories and their ownership and linkages. We have followed developments, some of which have been occurring during the writing of this report, but, inevitably, we offer only a snapshot of market technology which is likely to be superseded in some of its details in a matter of months if not weeks. Nevertheless, we are confident that many of our central contentions and recommendations will continue to have relevance.

The authors would like to thank virt-x Exchange Limited for sponsoring this study, and to express gratitude to all the interviewees and reviewers, particularly Robert Macrae, Pete Harris, James Hartley and Fritz Klein. The opinions expressed, and the conclusions drawn, are the independent judgements of the authors and do not necessarily conform to those of virt-x or any other party.

17 December, 2002

Keywords:

Best execution, cross-border trading, European equities, price benchmark, European superbenchmark, European market integration, exchange interfaces, clearing and settlement integration, order routing, XML standards, ISO 15022, MDDL, ICSD, CSD.

Table of Contents

Introduction

Best execution, broadly understood as the fiduciary responsibility of brokers to execute in the interests of their customers, is currently attracting a great deal of attention both in the US and Europe[1]. In particular, it has been seen as important to clarify and redefine the nature of a broker's responsibility to his clients, in the light of technological advances (e.g. order routing software) and the emergence of alternative execution venues to the major national exchanges. The UK Financial Services Authority (FSA) and the Committee of European Securities Regulators (CESR) have been exercised by this, and the FSA has recently published a consultation paper setting out its recommendations (FSA, 2002). The purpose of the present report is to explore the technological issues concerning the adoption of a pan-European best net price calculation as a benchmark for analysing and implementing best execution requirements.

At present no package is specifically marketed to provide automated routing to the European execution venue which offers the best price after deduction of all direct costs. Our main focus in this study is to identify the elements of technology necessary for this purpose (Section 1), to investigate the current state of availability of systems performing the three main component functions - pan-European price consolidation (Section 2), real-time net price calculation (Section 3) and "smart" routing to the best price (Section 4) - and to explore how these components could be brought together in an affordable way (Section 5). These purely technological issues touch on regulatory matters which we consider in the second half of the paper. We discuss there the suitability of specifying a pan-European price benchmark (Section 6), of mandating linkages between clearing and settlement systems (Section 7) and of standardising exchange interfaces (Section 8).

As a result of the research reported here, the authors are confident that all the technological elements of a pan-European best net price execution system are already available. Moreover, these elements can be packaged into very affordable products able to cope with (i) the large number of execution venues, (ii) the substantial amount of real-time data required to feed the netting calculations and (iii) the routing of orders in response to those calculations. We also believe strongly that this understanding of "best execution" would be in the interests of all investors (particularly retail investors) and should be adopted by the regulators. The report includes a number of other regulatory suggestions which, it is argued, would facilitate this focus on net price based best execution. These suggestions, in the areas of exchange interfaces and clearing and settlement linkages, should not be costly to implement, would lower costs to the industry and greatly facilitate moves towards both improving customer protection and providing a more integrated and efficient European equities marketplace.

1 Technological Definitions

This section defines the elements of technology, hardware and software, necessary for delivering execution at the venue offering the best net price. Later sections will look at how these elements may be used to provide the functional parts of a technological

[1] In our review of the non-technological literature, we have found Bacidore et al. (1999), Battalio et al. (2000), Harris (1996), and McCleskey (2001), (2002) particularly helpful.

solution, and how the components may be packaged so that they are affordable even to the smallest firms.

1.1 Hardware and Telecommunications

The necessary physical connections from the broker to multiple exchange venues, in one direction, and to the clients in the other, are represented in Fig 1.1

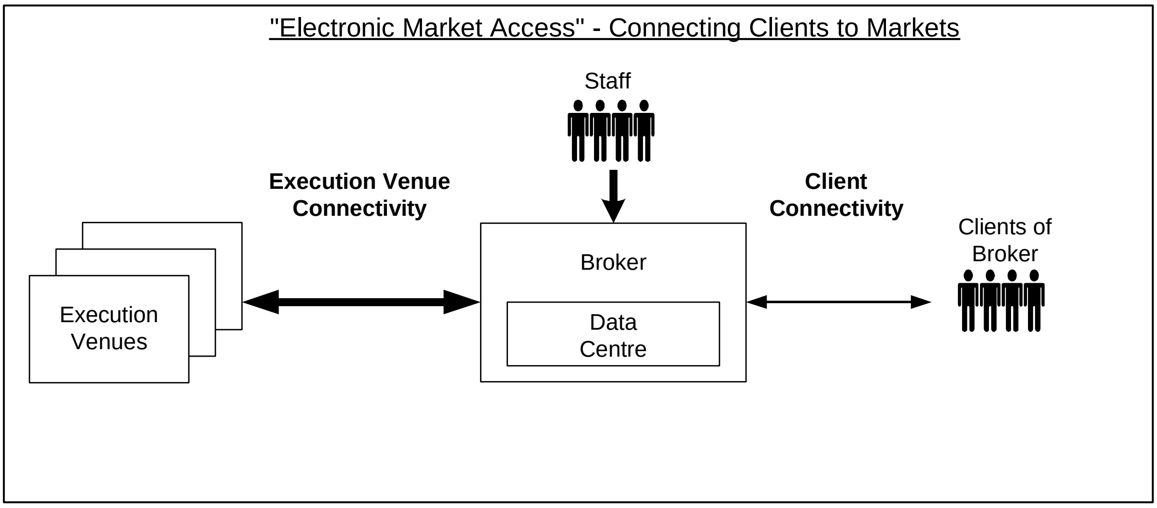

© Traderserve 2002

Fig 1.1

1.1.1 Physical Connectivity to Multiple Execution Venues

Pan-European best execution requires dedicated telecommunications links and back-ups to all the execution venues involved. For each venue, the broker must have links which handle both incoming data (e.g. order book[2], generic feed data, status information and fills for customer orders) and outgoing data (e.g. participating quotes, orders and requests for data). The links need to be fast and to offer sufficient peak bandwidth to cope with demand. The speed and reliability of order transmission to execution venues is of paramount importance.

1.1.2 Physical Connectivity to Clients

Historically brokers received orders from their clients by telephone, and their staff then phoned the orders to the execution venues. With the advent of electronic markets it is now commonplace for firms to provide electronic market access to their clients. Physical linkage to clients can involve either dedicated telecommunication links or virtual private network (VPN)/ Internet connectivity. The choice is dependent partly upon the client's requirements for guaranteed bandwidth, but an additional consideration, which will prove to be very important to our later discussions in Section 5, is the question of which parts of the system are to be deployed at the client's site. The bandwidth requirements are

[2] The "order book" for a stock, which is offered on an electronic exchange or for which there is an alternative market available on an ECN (Electronic Communications Network), is the set of limit orders current in the market. A "fill" occurs when an order received by the exchange matches a limit order. Normally, execution follows a price/time priority. There are many different kinds of market: Smant (2002) offers a useful classification.

critically dependent on the size and frequency of messaging, which itself depends on where the bulk of processing is being performed: on the client side or the server side.

1.2 Software: Applications and Modules

All of the necessary software modules and applications require hardware and networking to operate. Where large amounts of data have to be consolidated and processed quickly, fast CPUs and networking are advantageous, but for those parts of the system driven by human agent interaction rather than data, these are of less importance.

1.2.1 Logical Connectivity to Multiple Execution Venues

For simplicity we group logical, protocol, messaging, API (Application Programming Interface), vocabulary and data field definition issues together here and consider this as the "interface" to each execution venue. Virtually every execution venue in Europe and indeed globally has its own interface. This "tower of Babel" means that any Independent Software Vendor (ISV) or data consolidator wishing to provide exchange connectivity has to employ huge resources in interfacing their software to each execution venue, including processing data feeds, transmitting orders, receiving trade confirmations, exception handling and conformance testing.

The objective of most "execution venue connectivity consolidators" is to provide a single API or common interface to which other trading software modules can be connected.

1.2.2 Net Price Consolidation

There are no great conceptual difficulties involved in moving from simple gross price consolidation (of prices from multiple execution venues) to net price consolidation. All that is required is the deduction of the transaction costs incurred in making the trade. We do however make a distinction here between different kinds of consolidation:

1.2.2.1 Generic Net Price Consolidation

This requires access to a maintained source of execution venue and stock-specific market information such as costs and certain corporate actions.

1.2.2.2 Client/Order Specific Net Price Consolidation

This requires a Client Management System capable of maintaining client and order size specific cost information.

1.2.3 Order Routing

We distinguish four kinds of order routing:

1.2.3.1 Basic Order Routing

Rapid transmission of an order of a supported type via an API to a designated execution venue. This often utilises the common interface of execution venue connectivity software, and can be performed via either a centralised messaging core or a distributed system.

1.2.3.2 Generic Smart Order Routing

This involves automatic selection of the execution venue, and, where necessary, re-routing of orders which fail to execute. It is normally achieved by applying fixed rules. Conventional smart order routers could be adapted to handle multiple currencies and generic costs.

1.2.3.3 Client Specific Smart Order Routing

This is similar to the above but the order routing choice is in part influenced by client-specific details, including client-specific costs or order routing preferences.

1.2.3.4 Execution Strategies

We define Execution Strategy modules to be routers which offer an automated way of breaking up an order and of "working" component parts of the order on one or more execution venues. Modules differ in the extent to which they allow customisation. Some allow the incorporation of user-defined parameters, while the most flexible enable users/dealers to change the underlying algorithms without the use of programmers.

They are distinguished from all of the above categories by the support which they offer for *contingent* execution, in addition, or in place, of routing to the best price currently available.

Execution Strategies are vital when minimising market impact of larger institutional orders. This can be achieved by "slicing and dicing" larger trades into small limit or market orders. These orders are submitted to, and cancelled from, the market in response to a number of real-time considerations including price movement, width of spread, market depth information[3] and the passage of time. Randomisation, both in respect of the size of the order and the time of its submission, may also be used to disguise trading activity so that orders are not too easily "read" by the market.

1.2.4 User Order Management Software

A final category of software likely to be part of a complete automated best-execution solution is order management software, which allows dealers and/or clients to monitor and interact with orders. In larger participants this functionality will normally be provided by a commercial (or proprietary) Trade Order Management System (TOMS).

Detailed analysis of TOMS or other user order management tools is beyond the scope of this report. TOMS are selected on many criteria, often quite distinct from order routing and best execution. When TOMS are installed they are usually interfaced to existing or new order routing systems and linkages, as well as to back-office systems. It is expected that those TOMS vendors who already provide execution reports would add the supporting data for any new mandated price benchmark to orders recorded in their systems. Such an enhancement would be minor for most vendors.

[3] Market depth information is the record of "declared interest" in the market. For an order-driven market this is the limit order book itself; for a quote-driven market it is the complete set of bids and offers along with size and (in some cases) market maker.

This completes the introduction to the hardware, networking and software issues involved in net price based best execution, and the identification of the elements that need to be put together in a solution. Before we consider in Section 5 how such a solution might be make available at low cost, we consider the current state of technological progress concerning those elements.

2 Price Consolidation

Consolidation of trade, quote and limit order data from multiple execution venues for the same stock is the first of the necessary technological functions. We refer to the experience in the US to show how price consolidation there has been operating successfully for a quarter of a century, and how this is seen to have benefited the market generally, and the retail investor specifically. We then turn to the case of Europe to see what technological obstacles exist here to similar price consolidation.

2.1 The US Experience

2.1.1 The US National Market System

The US National Market System (NMS) was mandated in 1975 under Section 11A of the Securities Exchange Act (the '75 Amendments) to consolidate quotes from the major exchanges and regional exchanges into a National Best Bid and Offer (NBBO), and to consolidate trades with volumes onto a tape. Our description of the NMS is based principally on material from the SEC "Report of the Advisory Committee on Market Information: A Blueprint for Responsible Change" (SEC, 2001a).

The NMS has consisted of four separate parts or "plans", one dealing with options (the Options Price Reporting Authority plan - OPRA), two dealing with listed equities and one with Over The Counter (OTC) equities. The three plans for equities are:

- Network A covers all New York Stock Exchange (NYSE) listed stocks and consolidates their prices with those available on the regional exchanges: Boston Stock Exchange, Chicago Stock Exchange, Cincinnati Stock Exchange, Pacific Stock Exchange and Philadelphia Stock Exchange.
- Network B covers all listed stocks not on the NYSE that are either listed on Amex or are listed on regional exchanges and meet Amex listing requirements.
- Nasdaq/Unlisted Trading Privilege Plan (UTP) covers Nasdaq National Market securities. Nasdaq prices are consolidated with those on the regional exchanges. Recent changes to the UTP plan mean that Electronic Communications Networks (ECNs), offering electronic markets in Nasdaq securities, now contribute their best quotes to the consolidated quotation feed via SuperMontage, the Alternative Display Facility (ADF)[4], or by reporting to a local exchange (SEC, 2002).

[4] The ADF was introduced by the SEC as a response to concerns expressed by ECNs and competing exchanges who saw SuperMontage as an anti-competitive threat to their businesses. ECNs can now contribute to the consolidated quotation feed, by way of the ADF, without being part of SuperMontage. See Selway (2002b) and SEC (2002) for more details.

Since the SEC brought in the “Display” Rule 11Ac1-2 in 1980 data vendors have been obliged to display the consolidated information to avoid favouring one execution venue over another.

Historically it has been possible to find better published prices than the NBBO at ECNs which were not contributing to the NMS consolidated feed. Even now, the NBBO may not represent the best prices available to traders, for it can only consolidate the *declared* interest in the securities in question; it cannot take account of market makers or specialists who have not declared their interest (as they are perfectly entitled to do) but who are, in fact, prepared to deal within the NBBO. For these reasons, the NBBO has in practice often been easy to beat. Its value has not consisted in its perfection as a benchmark.

2.1.2 The Value of the NMS

2.1.2.1 Investor Protection

The SEC views the display of consolidated quotes and trades as the “cornerstone” of the US NMS because of its role in facilitating best execution, promoting investor protection and mitigating market fragmentation (SEC, 2001a). In reviewing the vision behind the ’75 Amendments that led to the NMS, former SEC Chairman Arthur Levitt, expressed the point strongly:

> *It is a vision rooted not in orthodoxy, but rather, in a practical recognition that investors are best served when diverse markets – exchanges, dealers, and alternative markets – compete for business; a vision where the best prices in any market are visible and accessible to all; a vision that embraces the goals of competition, transparency, market connectivity, and best execution, but is mindful of the inherent tensions among them.*
>
> *The wisdom and coherence of the framework lies in a single-minded focus: protecting investors* (Levitt, 2001).

In other words, the NMS has many benefits but the principal reason for its existence is investor protection.

2.1.2.2 Best Execution Monitoring

Following the imposition of SEC Rules 11AC1-5 and 11AC1-6 (Rules 5 and 6), brokers and exchanges have been required, amongst other things, to produce statistical data on their execution quality in comparison to the NBBO. We direct the reader to the SEC paper “Final Rule: Disclosure of Order Execution and Routing Practices” (SEC, 2000) for full details. Here we note that Rule 5 requires market centres to record, for different categories of order size, the average “effective spread” where this is defined as

> *the share-weighted average of effective spreads for order executions calculated, for buy orders, as double the amount of difference between the execution price and the midpoint of the consolidated best bid and offer at the time of order receipt and, for sell orders, as double the amount of difference between the midpoint of the consolidated best bid and offer at the time of order receipt and the execution price* (SEC, 2000, “Text of Rules”).

Also, for each of the order size categories, they must record the number of executions occurring (1) at a price better than the prevailing best bid (or offer), (2) at the best bid (or offer) and (3) at a price outside the best bid and offer. So, although, "best execution" is not *defined* under SEC rules in terms of the consolidated price, the NMS is central to US monitoring requirements.

Rule 5 has not been universally well-received in the US, partly because of the additional burden in recording and reporting that it places on firms. The effect of this, which has been felt by small firms in particular, has been somewhat mitigated by the emergence of software and ASP (Application Service Provider) services, now in use by both large and small institutions, which consolidate and publish the required statistics. In any case, it is not our intention to argue here in favour of Rule 5 reporting in particular, although, as traders, we *do* find value in many of the reported numbers. We merely record our support for the general principle of benchmarking trades against price. Section 6 deals with this issue in a pan-European context.

Critics of this focus often complain that price is only one dimension of "best execution", but the SEC responds to this objection as follows:

> *Rule 11Ac1-5 is needed, not because price is the only important factor in routing orders, but because there currently is little or no public information that would allow investors to assess a broker-dealer's handling of its customer orders.* (SEC, 2000, III.A.1),

and it writes of the intention of Rule 5 being to remedy a "glaring absence of public information" (SEC, 2000, III.A.1). This absence is felt, we believe, even more forcibly in Europe than the US, and, historically, less has been done to address it.

2.1.2.3 Enforcement

The NBBO has also been a crucial point of reference when firms are charged with failure to fulfil their fiduciary duty of best execution for their clients. Two key cases, Newton v Merrill Lynch (USCA, 2001) and the Geman case (SEC, 2001b), have shown that best execution obligations may go beyond simple execution at the NBBO if the broker has reason to believe he could obtain a better price, but this only shows that the NBBO is not enough on its own, not that the benchmark is dispensable.

Another thing that emerges from scrutiny of the Geman case is the existence of services to brokers which offer guaranteed price improvement over the NBBO - the existence of these was material to the case. This is highly relevant to the dispute considered in Section 6 concerning the value of pan-European benchmarking, since these services, which improve the quality of execution available to investors, would not exist were it not for the NMS.

2.1.3 The Future of Consolidation in the US

The SEC Advisory Committee on Market Information has recommended moving from the existing single consolidator model to a multiple-competing consolidators model, but they are at pains to stress that this is not because the existing system is "broken"; rather it is because of their desire to promote competition in this area (SEC, 2001a). The SEC sees this new competitive model as introducing new technological issues, in particular, the sequencing of information between different consolidators and the use of different formats and protocols by the competitors, but, significantly for our discussions in Section 8, they do not view these objections as insuperable.

At the same time unofficial consolidation has graduated to market depth information - i.e. displayed interest outside the best bid or offer. The Archipelago "Super-ECN" has consolidated the order books of ECNs for a number of years, and both Nasdaq SuperMontage and the Alternative Display Facility now consolidate some market depth information from participating market makers and ECNs (Selway, 2002b). The consolidated best price is used as a basis for order routing as described in Section 4.

2.2 Pan-European Consolidation

There are already vendors in Europe who offer customers a consolidated best bid and best offer taken either from the national exchange or from virt-x (the cross-border exchange), and our interviews revealed more than one candidate for consolidating post-trade information, should a consolidated tape similar to the NMS model be mandated in Europe. At least one vendor has already consolidated post-trade information as a proof of concept.

2.2.1 Hardware Issues

As we observed in Section 1.1.1, physical exchange links need to be fast and of sufficient bandwidth to cope with peaks of demand. For consolidators, the *relative* speed of links is also of great importance. A fair comparison between venues requires, for example, that the competing quotes that the customer sees are as near as possible contemporaneous. The geographical remoteness of one execution venue in comparison with another can impact adversely on the required simultaneity, but more often, when there are difficulties of this sort, bottlenecks in the communications are the dominant factor. For instance, for members to submit quotes and orders and to receive fills, some major European equity exchanges now recommend 2Mb lines whereas at least one allows only 64K circuits. The latter exchange also allows the consolidation of multiple members' orders down a single line by a service provider, which can help reduce the costs of increasingly popular ASP order routing systems, but is often impractical due to transmission delays caused by the currently low bandwidth. Another exchange transmits such a high volume of data during auction periods that at least one member turns off its communications at these times to avoid transmission overloads which can persist for significant periods after the auction.

For exchange data consolidators a vital requirement of hardware is that it allow processing which is sufficiently fast to avoid queuing. Market data tends to be concentrated in bursts of activity, so processing capacity must be much higher than would be required if data arrived at more regular intervals.

One further point relating to exchange connectivity: some execution venues specify very substantial hardware, in addition to multiple high bandwidth dedicated telecommunications circuits, in order that their members can run the electronic connection. This further increases the commissioning costs of physical connectivity to execution venues.

2.2.2 Software Issues

Speed is a crucial feature of exchange connectivity software as it is frequently under heavy load from multiple feeds and order sources. As such, the software itself can produce a key bottleneck in addition to those caused by bandwidth limitations and communications routing. Beyond this there are two central problems that software must handle: coping with multiple exchange interfaces and sequencing the data.

2.2.2.1 Differences in Exchange Interfaces

Differences in interfaces from market to market (as outlined in the Technological Definitions) form a barrier to competition with the established data vendors, but standardisation might also be welcomed by the major vendors so that they could concentrate their resources on adding value elsewhere. At present there are very few pan-European real-time consolidators and a larger number of domestic competitors in the various countries. This is principally because of the enormous resources required to interface and update the linkages to the many disparate data formats, protocols, and data vocabularies, particularly when execution venues have a tendency to "improve" their interfaces regularly and with little notice. The result of this is that vendors' development resources need to be held available at all times, which further increases costs to the financial industry, prevents vendor competition, and means that only the largest institutions can afford to create the connectivity in-house.

We return to deal with the issue of exchange interfaces in a section of its own (Section 8) where we recommend regulatory action which could greatly mitigate the problems raised here.

2.2.2.2 Sequencing

One software issue that has to be addressed by price consolidators is the sequencing of market data. This is not as trivial as it might seem because some exchanges fail to provide sequence numbers. Any attempt to resolve the sequencing of data, in this case, by increasing the granularity of time-stamping to the sub-second level could introduce other difficulties as it would also increase the size of each data packet associated with a trade or quote, and this might very well have bandwidth consequences for major data vendors. Amongst other things, this reinforces the point that software and hardware considerations cannot be dealt with entirely in isolation from one another.

2.2.3 Currency Conversion

Where securities are priced in different currencies, a real-time conversion needs to be performed to make the prices comparable. A major data vendor already offers this as a value-added service to customers, including hedge funds that are attempting to profit from arbitraging price differences between different exchanges, and others are ready to compete. A number of institutions already provide automated foreign exchange hedging of currency exposures such as those arising from cross-border trading. At least one Retail Service Provider (RSP) has publicly announced its use of these services.

2.2.4 The Need for Net Price Consolidation in Europe

Although there are at present practical difficulties that affect the ability of small vendors to compete, it would appear that there are no overwhelming *technological* reasons why a European Best Bid and Offer might not be provided along the lines of the NBBO in the US. However, in itself such a *gross* price calculation would be of very limited value and would certainly be of far less use to traders than its equivalent is in the US. The reason for this is that the whole cost of trading in Europe is critically dependent not just on the execution venue but on the clearing and settlement arrangements which differ vastly in their cost to the trader. In the US, the existence of a single monopolistic clearing and settlement house, the Depository Trust and Clearing Corporation (DTCC), ensures that whatever execution route an investor chooses to take, the total direct cost of trading will be much the same; in consequence, the best *gross* price available to a US trader is a more reliable indicator of the best *net* price (which is what really matters to the trader). This is not so in Europe. In Section 7 we consider in some detail what can be - and is being - done to improve transparency and efficiency of clearing and settlement in Europe.

We recognize that the case for price consolidation along the US model would be more immediately compelling if gross price comparison were more indicative of the most advantageous choice of execution venue available to the trader. In the absence of this it is necessary to focus on direct comparison of *net* prices, and we consider this in the next section.

2.3 Conclusions

2.3.1 Technology already exists capable of performing pan-European pre- and post-trade consolidation of quotes, order books and trades.

2.3.2 At present, pan-European consolidation is limited to a small number of international vendors. This is principally because of the complexity and costs caused by the multiplicity of interfaces used by different exchanges Europe-wide.

2.3.3 Regulations, along the lines proposed later in this paper, would address this barrier to competition.

2.3.4 The US experience is compelling both in respect of the technological feasibility and the competitive and regulatory value of price consolidation.

2.3.5 In Europe, gross price comparisons are more misleading at present, so simple price consolidation is of limited value and prices *net* of direct trading costs must be considered.

3 Real-Time Net Price Calculation

A simple comparison between displayed bids and offers is not sufficient to identify the exchange most suitable for executing a customer's order. A fair comparison should take into account different costs which that order would be subject to if executed at each of the competing venues. In this section we discuss the costs which should be netted, and we introduce a sophistication of the notion of best net price which takes into account the size of the client's order and makes use of available market depth information.

3.1 Allowing for Costs

3.1.1 Direct Costs

Costs to be netted on a real-time basis from the bids and offers include commission, taxes, exchange transaction fees, ticket costs, clearing and settlement costs. Where clearing and settlement costs are unknown in advance of trading (for example, because it is not known whether transfer across multiple depositories is required) we suggest that a worst-case figure should be assumed. This will encourage greater cost transparency and should help bring down those costs in time as execution venues compete for business on a net price basis. Execution venues and CSDs (Central Securities Depositories) would be strongly incentivised to publish accurate menus of costs for different equities in order to secure the greatest order flow where they can be the most efficient provider.

A number of vendors already provide software to perform netting calculations of the sort envisaged here. The calculations themselves are not very complicated, and if provided on an ASP model could be offered very inexpensively. Integration with a client management system to handle personalised costs is more challenging but solutions are available now. We discuss affordable solutions in Section 5.

3.1.2 Indirect Costs

In addition to direct costs, which may be accurately anticipated, transactions are also subject to indirect costs such as market impact (when the trader's own activity moves the market against him) and implementation shortfall (when the market moves in the direction anticipated by the trade, but before the trade is complete). Such costs, which are extremely difficult to estimate, should *not* be part of the netting calculation. We are considering here only orders of a size equal to or smaller than the total *displayed interest* on the offer (for buys) or on the bid (for sells), and for trades of this size indirect costs are usually negligible. We return to this point in Section 6 in considering the distinction between retail and institutional trading.

3.1.3 Soft Commissions

One further issue concerning costs: we note that part of the commission paid by a fund manager may be effectively passed back by the broker via the provision of research and software systems, resulting in an arrangement whereby the fund's investors are, in effect, paying for the fund manager's research. Such "soft commission" agreements, though very common, introduce a lack of transparency into commercial arrangements, and blunt the incentive of the fund manager to seek the best fills. They are relevant therefore to the best execution issue: they are indeed direct costs to trading

Schwartz and Steil, in their penetrating article "Controlling Institutional Costs: We have met the enemy, and it is us" (Schwartz and Steil, 2002), suggest that soft commissioning is at "the heart of the underperformance problem" affecting the fund management business, and that as a result of it fund managers are "hardly passive victims of sell-side structures and practices" as they would prefer to be seen.

The Myners report, "Institutional Investment in the United Kingdom: A Review" (Myners, 2001) has examined this issue recently and has concluded that

> *client's interest would be better served if they required fund managers to absorb the cost of any commissions paid, treating these commissions as a cost to the business of fund management, as they surely are. Fund managers would of course seek to offset the additional cost through higher fees; this would be a matter for them to agree with their clients. Under this system, the incentives would be different. Institutional clients would see more clearly what they were actually paying to have their funds invested* (Myners, 2001, p11).

In the case of conventional long-only equity funds, which tend to have relatively low turnover, this could deal with "soft commissions" and their impact on best execution. For fast trading hedge funds[5], on the other hand, the size of the commission would exceed any reasonable management fee as a result of the much higher turnover. Hedge fund managers tend to pay very close attention to transaction costs. Arguably, they have more incentive than conventional fund managers to keep execution costs low, for they are rewarded out of a percentage of net trading profits. Consequently, their interests are better aligned with those of their investors, and proposals along the lines of the Myners report are less necessary. In the words of one hedge fund manager, "we have the performance fee to keep us honest".

3.2 Market Depth

Where the size of an order *exceeds* the volume on the best bid or offer but does *not exceed* the total displayed interest, a more sophisticated notion of best net bid and offer is appropriate which looks beneath the best quotes to sufficient depth to fill the order. This is the Volume Weighted Average Best Net Bid and Offer (VWABNBO) for the required size. To illustrate this consider the following example.

The table below shows the top of the order book for a stock quoted at two execution venues, EV1 and EV2. The total size displayed at a price is shown alongside the bid and offer. In this example the best (gross) offer is on EV1 at 77.30 and the best (gross) bid is shared between the venues at 77.25.

[5] A hedge fund is distinguished from a long-only (mutual fund) principally by the use of derivatives, other instruments or short-selling to offset market risk, and leverage to increase the size of returns (proportionately with risk). Hedge funds usually charge their investors two fees: a *management* fee which is a percentage of the value of assets managed (e.g. 1-2% annually), and a *performance* (or *incentive*) fee which is a percentage of net trading profits (e.g. 20% of quarter-end profits above the previous high-water mark of the fund). According to *vanhedge.com* global hedge fund assets stood at about $600 billion in December 2001

Gross Order Books

EV1			
Size	Bid	Offer	Size
10000	77.25	77.30	5000
10000	77.20	77.35	5000
5000	77.15	77.40	10000
10000	77.10	77.45	10000

EV2			
Size	Bid	Offer	Size
10000	77.25	77.35	20000
50000	77.20	77.40	40000
40000	77.15	77.45	20000
25000	77.10	77.50	15000

Suppose that the estimate of the cost of trading at EV1 for Customer1 is 3 cents per share and at EV2 is 2 cents per share. (We assume for this example that there is no dependence of these costs on the size of order). The effective net order book would then look to the customer like this:

Net Order Books for Customer1

EV1			
Size	Bid	Offer	Size
10000	77.22	77.33	5000
10000	77.17	77.38	5000
5000	77.12	77.43	10000
10000	77.07	77.48	10000

EV2			
Size	Bid	Offer	Size
10000	77.23	77.37	20000
50000	77.18	77.42	40000
40000	77.13	77.47	20000
25000	77.08	77.52	15000

3 cents has been subtracted from the bids at EV1 and added to the offers, and 2 cents has been netted at EV2. As a result of this, EV1 has the best net offer at 77.33 and EV2 has the best net bid at 77.23. If the customer wishes to buy up to 5000 shares he will expect to pay an effective net price of 77.33 by executing at EV1, which is better than he can achieve at EV2. Suppose, however, that he wishes to buy 20000 shares instantly. At EV2 he can expect to pay a net price of 77.37 because the (net) order book shows that the entire size of order is available at that price. At EV1, on the other hand, the best he could expect is 5000 at 77.33, 5000 at 77.38 and 10000 at 77.43, which would give a Volume Weighted Net Offer of

(5000 * 77.33 + 5000 * 77.38 + 10000 * 77.43)/(5000 + 5000 + 10000) = 77.3925.

This price compares unfavourably with the 77.37 offered at EV2, so for this larger order size the Volume Weighted Best Net Offer has moved from EV1 to EV2.

This example shows the importance of taking into account the size of the order when choosing the execution venue. It is not appropriate to send the order in all circumstances to the venue with the *best* net bid (or offer). The *depth* of the order book should also be examined when the size of the order exceeds the volume shown at the best bid (or offer).

One vendor already offers for the London market a display showing the Volume Weighted Price on the Bid and the Offer at each price level, assuming all volume is “swept” to that depth. They intend to extend this to the entire European market. Admittedly, this is not currently combined with netting calculations, but it could be made so.

Instead of identifying the best single venue for execution, an order may be split across multiple execution venues to take advantage of *all* the venues offering the best price. This requires smart order routing of greater sophistication than the basic on offer (see Section 4), and is a more viable option in the US than in Europe where the fragmentary nature of clearing and settlement imposes additional costs that normally mitigate the benefit of splitting the order between multiple venues. We consider clearing and settlement in some detail in Section 7.

3.3 Conclusions

3.3.1 Technology already exists to perform the calculations necessary to identify in real-time the best available price net of all explicit costs (including commission, taxes, clearing, settlement) and allowing for currency conversion where necessary.

3.3.2 A Volume Weighted Average version of the best net bid and offer (VWABNBO) requires a little more computing power, but the calculations are straightforward and are already in use. This is suitable for orders of a size greater than that available at the best bid (or offer), but no greater than the *total* displayed interest on the bid (or offer). We strongly advocate its use in the calculation of a European super benchmark

3.3.3 Where clearing and settlement costs are unknown in advance of trading, we propose that a worst-case figure should be assumed, which should encourage greater transparency and lower costs as venues compete for business on a net price basis.

4 Order Routing

The purpose of smart (or intelligent) order routing algorithms is to obtain for the trader the best price available, across multiple execution venues, for his order. We discuss here the features of these algorithms, their use in the US equity markets and issues concerning their adoption in a pan-European context.

4.1 The US Experience

While many details of smart order routing algorithms in use in the US remain proprietary, there are generic features. A typical such algorithm used by a broker is described below.

Smart Order Routing Algorithm

Step 1 Receipt of a limit or market order
If the order is a market order proceed to Step 2.
If it is marketable (i.e. it is a buy order at a price >= the National Best Offer, or a sell order at a price <= the National Best Bid) proceed to Step 2.
Otherwise (in the case where the order is a limit order at a price more optimistic than the current NBBO) route the order to the NYSE for listed stocks or a particular ECN for OTC stocks. Until filled, periodically check whether the order has become marketable. If so cancel the order and proceed to Step2.

Cont.

Step 2 Cap the order to the NBBO range
If the order is not "tight enough" to the market – e.g. it is a limit order to buy at 77.35 when the best offer is currently 77.32 – cap the order to the NBBO range (e.g. put the limit order in at the best offer 77.32, though other algorithms can be used). Then proceed to Step 3.

Step 3 Route the order to the best price
Identify the execution venue offering the best price and send the order there. In the case of a tie between two or more execution venues, one is chosen on the basis of some ordering which may take into account, for example, the expected speed of execution. Proceed to Step 4.

Step 4 Check state of execution
If the order executes fully, the transaction is complete.
Otherwise check whether the order is still marketable at the chosen execution venue.
If it is proceed to Step 5.
Otherwise proceed to Step 6.

Step 5 Check passage of time since execution attempt started at this venue
If 10 seconds (say) have elapsed then proceed to Step 6.
Otherwise proceed to Step 7

Step 6 Check other venues
If the order is marketable at an alternative venue proceed to Step 8.
Otherwise proceed to Step 4.

Step 7 Check the price of stock at all venues
If there is a better price, proceed to Step 8.
Otherwise, return to Step 4.

Step 8 Restart routing decision process
Cancel the balance of the order and return to Step 2 with a new order for the balance.

End of Algorithm

It is significant that the NBBO benchmark is an explicit part of this algorithm - a fact which underscores the value of the US benchmark for assisting execution that is in the interests of the retail trader.

From the above description it is also plain that smart routing may involve more than simple identification and routing. A great deal of complexity, for example, in handling failed cancellations, is omitted from this relatively high-level description of the algorithm, but even so, the algorithm is designed to take advantage of changes in the NBBO while execution is incomplete and to cope with difficulties in obtaining fills from the chosen venue.

There are many other sophistications employed by "smart routers". Some algorithms split the orders across multiple execution venues, allowing the trader to benefit from all sources of the best price when the displayed interest is not great enough at a single venue to fill the required order; some allow the choice of venue to be influenced by an order routing priority list; some take transaction costs into account, while others involve the collection and use of statistics concerning the speed and/or historical success-rate of execution at the various venues. Statistics-based routing of this sort goes beyond simple best execution: these algorithms may be used to offer price improvement over the NBBO, but, as previously argued in Section 2.1.2.3, the existence of the benchmark is a precondition for the existence of such services and the ability to analyse their value.

More sophisticated still are the routers that hunt out *undeclared* interest at an execution venue, by initially routing a small part of the order and then monitoring whether further volume appears at that price. These are on the border of the distinction which we draw between "execution strategies" and "smart routers", the latter of which we take to be strategies designed to offer the trader the best price *now* for a market or marketable limit order.

Some algorithms purporting to be "smart routers" actually build in preferences for one venue over another on grounds other than that of finding the best price for the trader. For example, it is common for an ECN to sweep its own book before passing the remains of an order on to its competitors. This is clearly *not* best execution.

In all, it is estimated by transaction cost analyst experts, whom we have interviewed in connection with this study, that between 10 and 20 per cent of trades in the US are subject to smart order routing.

Nasdaq's SuperMontage is, in part, a response to the growing popularity of this more intelligent approach to trading given multiple pools of liquidity. SuperMontage takes and displays, in a single consolidated view, quotes and limit orders from participating market makers and ECNs, and it offers three algorithms for routing orders to the participants: (1) Price-time priority like a classical limit order book, (2) price-size priority allowing a trader to reduce the number of partial orders by effectively sorting the quotes/limit orders occurring at a single price on the basis of size rather than time, and (3) price-time with fee consideration priority which would route orders to ECNs charging a fee only as a last resort.

SuperMontage doesn't offer a perfect example of smart order routing because, in addition to the fact that its routing algorithm is quite rudimentary in comparison to the best alternatives, it is limited in the destinations to which it can route: an ECN that chooses not to participate in SuperMontage may well have a better price, but this will be ignored. At the present time it seems very unlikely that all ECNs will choose to display their liquidity on SuperMontage in the long term. The major ECNs by share volume are the newly merged Instinet group consisting of Instinet and Island, and Archipelago which has recently completed its merger with REDIBook[6]. Island has moved its prices to the Cincinnati Exchange while Instinet uses the ADF, and Archipelago is converting itself from an ECN to an open electronic exchange (ArcaEx) - the equity trading arm of the Pacific Stock Exchange. When the roll out of Nasdaq stocks on ArcaEx is complete,

[6] On recent figures the Instinet group has approximately 30% of share volume in Nasdaq shares and Archipelago about another 20%. Both are now targeting the NYSE where their current volumes are estimated at <5% and 2-3%, respectively. Figures are based on reported volumes in October and November 2002 on the Instinet website *http://www.instinet.com/trade_data/trade_data_month.shtml* and on the Archipelago website *http://www.tradearca.com/data/Volume/monthly.asp.*

which is expected some time in the first quarter of 2003, Archipelago will stop posting its best bid and offer to SuperMontage.

US investors are offered a wide variety of smart order routing solutions and clearly find them beneficial and affordable.

4.2 Pan-European Order Routing

Order routing services can be offered only by appropriately authorised brokers. They have traditionally deployed execution venue connectivity software, much of which is currently provided by ISVs, for their staff to enter orders manually into electronic markets. However, nowadays, many brokers also offer direct market access (electronic order entry) to their clients. This is usually achieved by additional software modules beyond the basic connectivity system and may be combined with other service delivery to the client. Smaller brokers can outsource all or part of the electronic connectivity to exchanges and to clients. Some broker business models are shown in Fig 4.1.

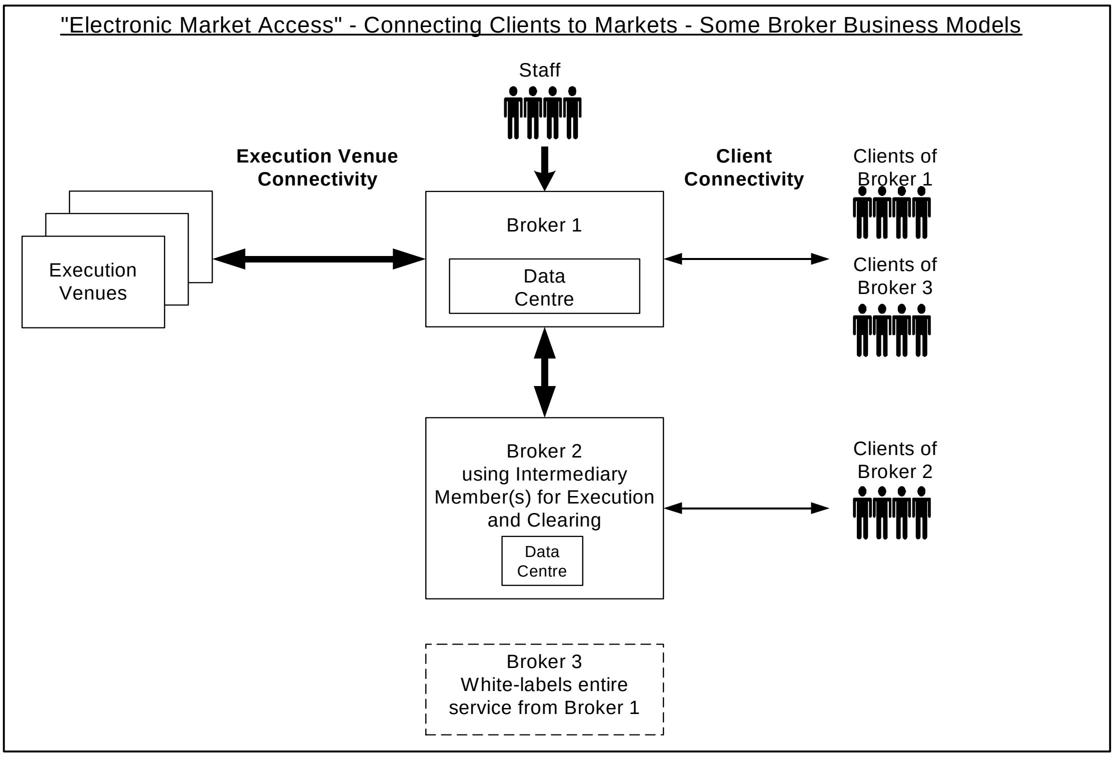

© Traderserve 2002

Fig 4.1

A number of ISVs, including one whose systems carry up to 80% of Xetra volume, have the relevant connectivity and trade order management software to develop best net price order routing in Europe at reasonable cost. Various existing software modules contain all the required elements of technology. Another vendor already offers value-added modules to allow exploitation of price differences between European exchanges.

Concerning the choice of medium for providing physical connectivity to clients, we note that the internet can be a very effective low-cost solution for delivery of small packets of information in many circumstances, but is much less suitable for delivering real-time high bandwidth data feeds between computers. When, for example, small pieces of information are sent infrequently to a responding human, communications delays tend to be small in comparison with the human reaction time.

In the context of best execution, if the calculations to determine best net price are performed at the client site, then the limited and variable internet bandwidth is likely to be inadequate to deal with the real-time market depth data from multiple exchanges. If, on the other hand, those calculations are performed at the broker's site, using customer-specific information where necessary, the customer needs very much less data, and internet connectivity may suffice. The use of net price, and the need to deal with client-specific costs and order-size-specific volume weighted averages, introduces these architectural complications which are not needed in conventional US smart-routing solutions. However, such a service could be installed locally in the broker's data centre or in an ASP centre, centralising connectivity, hardware, software and support, which should make it very affordable for smaller firms. The next section considers ASP solutions in more detail.

Clearly, it would not be reasonable to expect every firm to be a member of all exchanges, but we suggest that intermediaries and clearing members would continue to be used as at present by small firms (see Fig 4.1). This can simplify linkages and reduces costs.

Finally, we draw attention to the fragmented nature of clearing/settlement systems in Europe. Linkage between them all is a clear pre-requisite for the viability of *universal* pan-European smart routing across *all* execution venues, for, at the moment, the costs, delays and risks involved in cross-border transactions between unlinked systems frequently remove the benefit of finding a better price. We defer this issue until Section 7 where we show that it is being addressed *now*.

4.3 Conclusions

4.3.1 Technology already exists to route orders to the European execution venue offering the best net price.

4.3.2 In the US "smart routing" is heavily used by retail investors, and is found beneficial and affordable.

4.3.3 For Europe, there are additional issues with universal smart routing such as membership requirements and the current partial linkages between separate clearing and settlement systems, but these are all avoidable or in the process of being solved.

5 Making the Technology Affordable

Although all the technology elements to achieve pan-European net price based best execution exist, they are not currently marketed as a package for this purpose. In this section, we suggest configurations in which the technological elements necessary can be made available to even the smallest participants on a cost-effective basis.

5.1 Possible Configurations and Cost/Benefit considerations

We shall consider the costs implicit in deploying best net-price execution technology in three ways. The most important cost factors to consider are:

- Connectivity costs - variable in size and best managed by collocation of subsystems requiring connectivity which is fast, high throughput and multiply redundant.
- Support costs - usually large, minimised both by collocation and by the use of thin client technologies (e.g. browser-based) avoiding local configuration, installation and support at client sites.
- Price of processing power - usually negligible compared with the above, even when client-specific calculations using their individual direct costs are required, and particularly where distributed processing across inexpensive computer arrays is used (e.g. rack-mounted PCs).

Solutions must avoid locating the netting calculations (client-specific order routing) at a client's site remote from the main data servers and exchange connectivity (basic order routing). Such a configuration would cause processing to be critically constrained by the client connection and, in consequence, would likely result in reduced overall speed of processing and routing, as well as lowered reliability and increased costs for connectivity and support. All the configurations below avoid this constraint.

5.1.1 Private Data Centre

This configuration involves an institution deploying all the server technology in its data centres and supporting in-house users and remote clients. Cost considerations include:

- Physical and logical connectivity to all execution venues,
- Connection and maintenance of market information (e.g. corporate actions, ISINs etc.),
- STP (Straight Through Processing) links for processing fills,
- Incorporation of client- and order-specific costs,
- Client connectivity and support,
- Internal staff connectivity and support.

Exchange connectivity requires large numbers of private circuit links, considerable hardware together with a combination of commercial and proprietary execution venue consolidation systems from multiple vendors, and significant ongoing support of these linkages. The market information will normally be available from existing corporate systems, although, perhaps, subject to an internal charge. STP links will probably be already in place. On-the-fly extraction of client-and order-specific costs from existing databases to feed the net price calculation may require a significant systems linkage and integration effort. Considerable resources may also be required to provide and support client connectivity whether clients are connected by private circuit or are using suitable security over the internet (VPN or other). Finally, we expect that internal staff connectivity will be provided by existing Local Area Network (LAN) connectivity and that support will be provided by current technical support staff.

Clearly a solution of this type is appropriate only for large institutions. The additional cost of extending their architecture to include net price based best execution should not be large relative to their existing costs. This solution places all components under the direct control of the institution, but it is worth noting that even some large institutions are already moving to third party ASP provision (see Section 5.1.3). This is seen as a considerably more cost-effective solution to their client order-routing and servicing needs.

5.1.2 Solutions Using Intermediaries

In addition to the costs incurred in commissioning, running and supporting multiple execution venue links, there are also membership costs, the most expensive of which currently tend to be setting up and maintaining clearing and settlement - which in Europe, at the moment, can require local staff in many countries. It is probable, therefore, that many firms will continue to use specialist clearing members in different exchanges.

Smaller institutions at present often choose to route their executions through these intermediaries, who may also provide white label connectivity to the smaller institution's customers. This can remove most of the complexity and set-up costs of electronic market access for the smaller institution.

5.1.3 Third Party ASP (Minimal Cost Solution)

The logical endpoint of this progression is an ASP centre where a third party concentrates all physical and logical exchange connectivity and also customer connectivity and support. Such ASP models are starting to achieve widespread uptake in a number of arenas include equities trading and foreign exchange (FX). At least one of the largest international investment banks, and many smaller institutions, are already using third-party ASPs to deploy TOMS (Trade Order Management Systems) and execution connectivity to their equities customers. In FX, most of the top 100 global banks are using ASP systems for at least part of their client business.

A "web-native ASP" configuration (requiring only browsers and with no dedicated telecommunication links to clients) is shown in Fig 5.1. This example is based on one configuration of our own company's software architecture (with third party sub-systems and linkages), which we use solely for the purpose of illustration, emphasising that many vendors could quickly package the minimum set of elements required for pan-European net-price based best execution. We stress that we are not suggesting that any individual technology be mandated - firms should be best served by competing vendors and integrators with differing best execution technologies.

In such web-native configurations, all end-user access is via web browsers, which reduces the dependence on the end-user's computer configuration and allows the users to "roam" to different computers if required. It should be noted that some ASP operators prefer to use technologies such as WebStart which they say will download and launch full Java applications and thus reduce browser dependence. The authors are currently satisfied with applet deployment to browsers, as, in their experience, it has proved reliable over a reasonably extended period, has enabled simple upgrades and avoided many of the security, compatibility and support issues associated with local applications.

A Functional Architecture for Low Cost Pan-European Best Net-Price Equity Execution

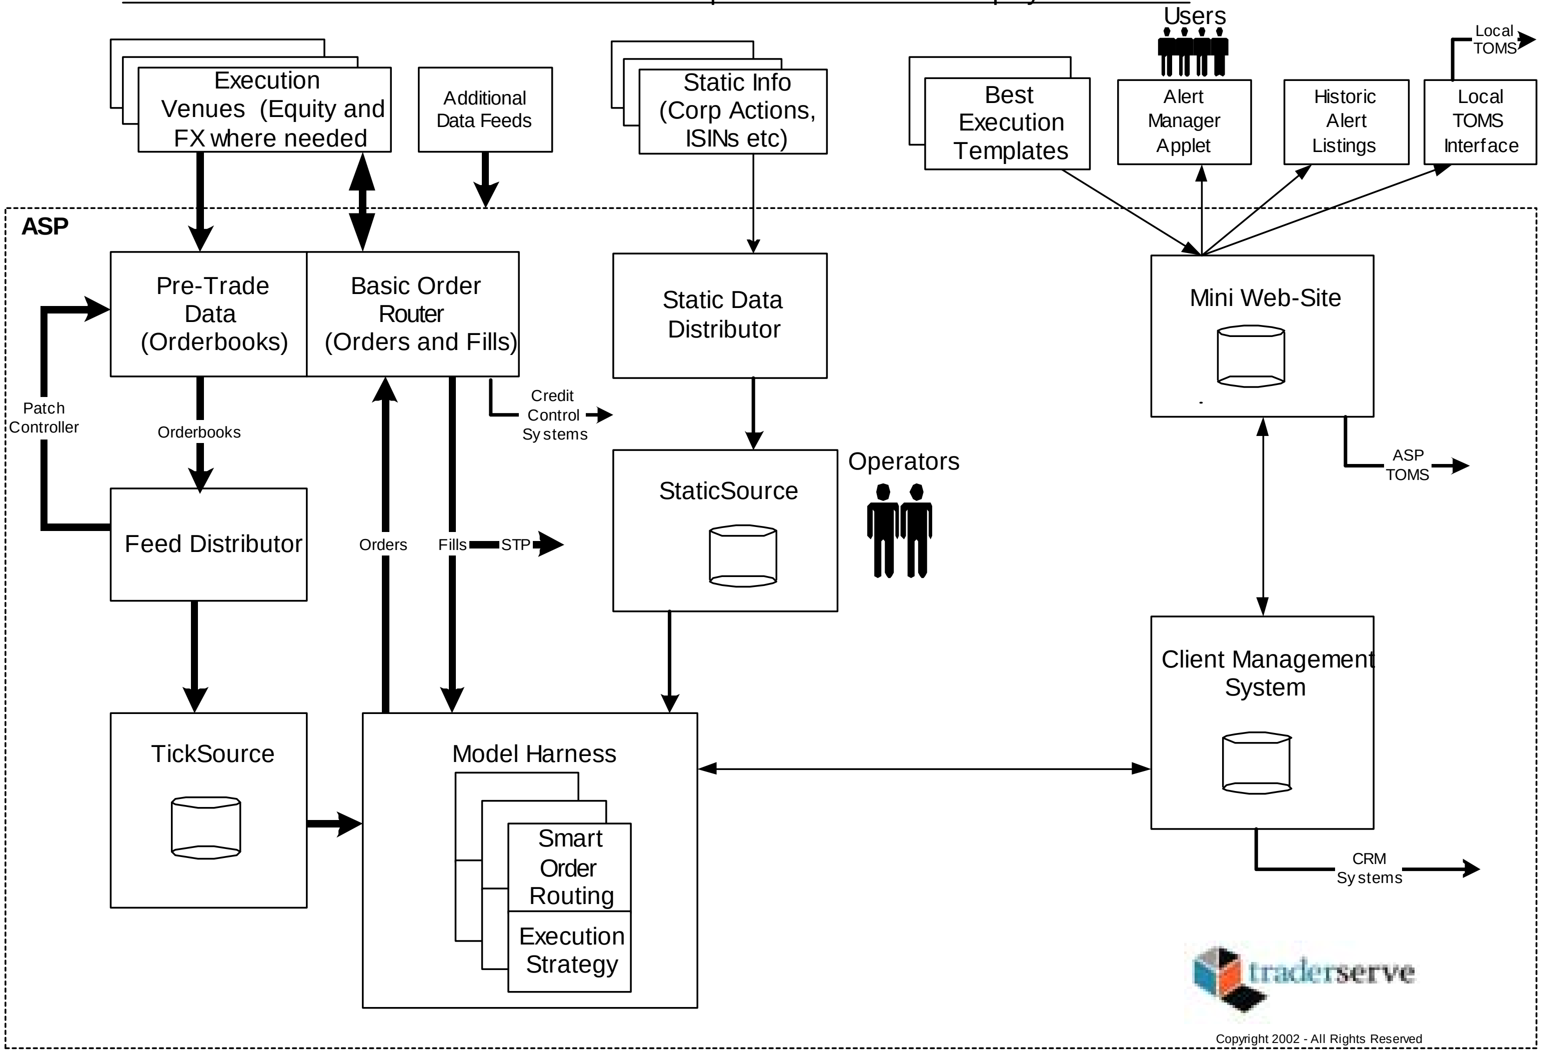

Fig 5.1

In the figure shown, the user (either in-house or remote, and either a staff member or customer), accesses the best execution system via a simple web page (template) or a link to a third-party TOMS. The template or link allows the user to enter the equity and quantity required and any additional preferences that its creator has allowed. An XML (eXtensible Markup Language) model, in this case an execution strategy or smart router, is then created in the user's private library in the client management system which also provides any client-specific information needed. The model loads into the model harness, where it can either be tested on historical data to evaluate the strategy or operated in real-time. In the latter mode it performs whatever smart routing is required and keeps the user informed by alerts (which will fall-back to SMS or email in the event that the user's connection is broken). Historical alert listings can be called up by the user, and it is also possible for the institution to administer its users remotely.

With our technology, the XML model is not hardcoded but is created by a user (usually a domain expert) in a browser-based graphical model-building environment. This offers complete flexibility for specifying real-time calculations and logic, and allows storage and re-use of a user's own XML constructs, greatly speeding development. Models can create any combination of alerts, orders, and data feed outputs, thus allowing a domain-expert rapidly to specify, modify and deploy his own smart order routing algorithms and execution strategies without the aid of a programmer. When the model is complete, it is converted to a template which exposes only the desired parameters to ordinary users.

We do not claim that an ASP solution is a universal panacea. There are still interfacing costs - for instance, the fills will still have to be routed to the institution's back-office, connections made to the credit control systems and the client management system linked to the appropriate customer relationship management systems (CRMs). Nevertheless, the advantages to small firms are evident. Connectivity and support issues are largely outsourced, creating very considerable savings with no reduction in flexibility. Performance and reliability can also be improved by housing in a centre with specialised equipment, communications and support - indeed small firms may benefit by placing other linked systems within the ASP centre.

Given that ASPs quite typically operate on a low maintenance fee plus a resource usage charge, the costs for small institutions could be very low indeed - low enough that they should form no barrier to a requirement from the regulators that even the smallest firms should provide net price based best execution.

5.2 Conclusions

5.2.1 A service combining best net price calculation on multiple markets with routing to the market offering the best price could be offered very inexpensively on an ASP model.

5.2.2 Costs should be sufficiently low that they would not present a barrier to making net price based best execution a regulatory requirement.

6 Price Benchmarking

Objections to operating a price benchmark are evaluated in this section. The value to investors of the US benchmark is amplified and the proposed VWABNBO pan-European benchmark (Section 3.2) defended. The restrictions on its use are clarified in the context of the distinction between retail and institutional trading.

6.1 The Objections to Pan-European Benchmarking

Sections 2, 3, 4 and 5 have shown that there are no overwhelming *technological* reasons preventing the adoption of a pan-European best net price benchmark, VWABNBO , for analysing and implementing best execution of European equities, and we have shown that, in the presence of a demand for the required technology, it could be provided very inexpensively and well within the budget of small firms. However, in the recent consultation document, the FSA is not recommending that a "super benchmark" European Best Bid and Offer be created by regulation citing "difficult decisions about consolidation of price information and market access" (FSA, 2002). Though the FSA document does not itemise any of these difficulties, it is reasonable to suppose that they lie in the areas which we have identified - namely, the requirement of access to multiple exchanges, the existence of multiple and imperfectly interconnected clearing and settlement systems and the need to cope with multiple different exchange interfaces. We answer these difficulties in Sections 4, 7 and 8 respectively of this report.

A further objection to setting a price benchmark, which we mentioned in Section 2.1.2.2, is that price is only one dimension of best execution amongst many (including timeliness and liquidity enhancement) and that the focus on price can be harmful if it causes the other dimensions to be overlooked. But we believe, and the SEC appears to concur, that price comparison, even when imperfect, is an important element in consumer protection. Where other aspects of execution are more important to the investor he might choose to waive his right to best net price execution or simply make use of other order types such as limit and not-held.

In any case, where measurable these other aspects of execution could be monitored alongside the price measures, such as effective spread for different sized orders, as they are in the US by software packages and ASP services which have emerged in response to the SEC Rules 5 and 6.

In Section 2.1.2.3 we referred to the Newton v Merrill Lynch and Geman cases in the US. These show the value of the NBBO benchmark in enforcing best execution responsibilities, and in lowering the cost of execution by giving rise to guaranteed price improvement plans. It does not matter, to these ends, that the NBBO is, as we have observed, an imperfect benchmark.

The authors can see no way in which best execution could be objectively judged or enforced without some price benchmark. This is of particular concern given the increasing prevalence of "internalisation"[7] of order flow within firms, where the absence

[7] We define "internalisation" as the process by which an institution, presented with a customer order, bypasses the regulated market, by matching the order with another customer order or against the trading

of a benchmark would make it impossible to decide whether or not certain customers are being favoured at the expense of others, with whom their orders are crossed. Post-trade reporting alone cannot remedy this.

6.2 Distinguishing Retail and Institutional Trading

Recently a *potentially* more substantial problem for the net price benchmark (VWABNBO) has been raised. This is the concern that exclusive attention to the direct (or explicit) costs of trading (such as commission and spread) at the expense of indirect (or implicit) costs (such as market impact and implementation shortfall) overlooks the major sources of transaction costs. In one study (Alba, 2002), from which the FSA document reproduces figures, it is claimed that direct costs comprise only about one third of the total.

These observations seem somewhat at odds with the FSA's main reason for rejecting the Stock Exchange Electronic Trading System (SETS) best bid and offer "safe harbour" for best execution - namely, that it is a standard that is too easily bested since, on their own figures (FSA, 2002), some 22% of all SETS trading occurs within the best displayed bid and offer. Indeed, where near instantaneous execution occurs in an electronic market, there can be *no* market impact for that order. Moreover, similar considerations would presumably apply to US equities trading where, as we have seen, a price benchmark is vigorously defended by the SEC and price improvement is offered as a standard service.

Despite this, the FSA consultation paper, in its evaluation of the case for a benchmark, places a great deal of weight on the existence of indirect costs. It is important then to stress that we are proposing a benchmark only for trades of a size *related to the current total displayed interest in the market*. Typically these are retail orders, rather than the huge orders common in institutional investing, particularly in portfolio transitioning. In the study mentioned above (Alba, 2002) the figures are drawn from a sample of trades for which the weighted average daily volume of trading represents 8% of the total volume for that stock on the day - this is very far from the conditions affecting private investors. Furthermore the VWABNBO benchmark is not for orders "sliced and diced" and worked into the market over periods of hours or days as is normal practice in the institutional world, but *market* orders to be executed on receipt or *marketable limit*[8] orders. Where *small* order slices from large institutional orders become marketable, they would also benefit from the protection offered by the benchmark. However, any attempt to define best execution which attempts to cater for both the lengthy process of filling huge institutional orders and the immediate execution of small market orders - as suggested in the FSA consultation document[9] - will struggle to cope with the complexity of executing large institutional orders (where indirect costs are key) while missing the opportunity to impose a clear and enforceable standard for the execution of small market orders.

---

position of the firm itself.. Some market participants believe that the investor will need very careful protection where internalisation is used. We refer the reader to Euronext (2002a) and (2002b), where they include "best execution" as a key issue.

[8] A limit order to buy at Price1 becomes marketable when the best offer becomes less than Price1; a limit order to sell at Price1 becomes marketable when the best bid becomes greater than Price1.

[9] "We also believe that the best execution obligation can be restructured so that it is flexible enough to cope better with the differing needs of both institutional and retail investors". (FSA, 2002, p13, 3.13).

The FSA consultation paper directs another argument against retaining a "best price" standard for retail orders which would count against our suggestion if it were cogent. A paraphrase of the argument is as follows: 'institutional customers of investment managers (e.g. pension funds) also represent the interests of many private individuals, so there would be an asymmetry in the duty of care being offered to those individuals whose investments are managed by pension funds, etc. and those that traded for themselves.' But the point is that the *nature* of the trading is so very different: a huge trading process on the one hand, a small market order perhaps on the other. The protection offered by the price benchmark is of great value even though its application is limited by the size of the total displayed interest.

One final point concerning the Alba paper: the conclusion that two thirds of the total costs of equities trading comes from indirect costs does not accord with this report's authors' extensive experience of operating very fast trading strategies in a number of liquid markets including Equity Market Neutral Trading[10]. Any order exceeding the total displayed interest in a stock has to be subject to some execution strategy (automated or manual), but for each *individual small slice* or relatively small orders, market impact costs are minimal due to their near instantaneous execution. One of the marks of a good execution strategy is then its ability to minimise market impact from one order slice to the next. Short term trading (including short term market neutral arbitrage) which has been a growing sector of the hedge fund management industry, could not exist if the indirect costs of trading using sensible execution strategies for reasonable order sizes were in general anything like as high as Alba indicates.

Clearly there are practical limits even with the best execution strategies. If the total size of the order slices executed across a period is required to be a very significant percentage of the total liquidity, the net result of executing all the slices is likely to be significant market impact. One specialist hedge fund manager which often trades 30% of the daily volume of smaller stocks over consecutive days estimate their indirect costs at 80% which is considerably higher than Alba's figure. This extreme case starkly illustrates the difference between the lengthy and complex institutional trading process and instant execution of small retail market orders. In the former case the order slices may not even be small enough relative to the total displayed interest to allow VWABNBO to be invoked at all. However, in the spectrum between this extreme case and retail trading, much institutional trading in liquid stocks in electronic markets can be efficiently done via execution strategies which slice orders; and where these slices are smaller than the total displayed interest they could also benefit from the protection offered by the benchmark.

6.3 Conclusions

6.3.1 The pan-European best net price is a suitable benchmark for European best execution of market - or marketable limit - orders of a size related to the total displayed interest in the market, and should be mandated.

6.3.2 More precisely, our suggestion is that brokers be required to direct customers' market orders (of a size not exceeding the total displayed interest) to the execution venue offering the VWABNBO price to that size, or,

[10] The figure of two thirds is also *twice* that calculated in an extensive global survey of equity trading costs by Domowitz et al. (2001) with data supplied by Elkins/McSherry Inc on 135 institutional trading firms. This suggests that the conditions of the Alba sample are exteme even for *institutional* trading.

alternatively, offer execution at a price at least as good. Sufficient contemporaneous data should be recorded to allow monitoring of compliance with this obligation.

6.3.3 This benchmark would offer valuable consumer protection for the *small market orders* typical to private investors. It would also assist price/cost transparency and help promote European market integration.

6.3.4 With increasing European moves towards order "internalisation" within firms, the lack of a price benchmark would make it impossible to counter the suggestion that favoured customers are receiving better execution at the expense of others.

6.3.5 The existence of a price benchmark is a precondition for the emergence of price improvement services, common in the US, which result in competitive improvement of investor's executions.

6.3.6 Best execution of retail orders could not be objectively judged or enforced effectively without a *price* benchmark.

6.3.7 Institutional trading of large orders is a different world where strategies of execution minimising market impact are commonly employed. While it makes sense to think that some of these strategies may be better than others, a clear concept of "best execution" for this *lengthy* process similar in clarity to our suggestion covering small orders for *immediate* execution is certain to be somewhat problematic.

6.3.8 Market impact costs, though they may be very high in the context of the execution of large institutional orders, are zero in the case of near instantaneous execution of small orders in electronic markets.

6.3.9 Where *small* order slices from large institutional orders become marketable they would also benefit from the protection offered by the VWABNBO benchmark.

7 Cross-Border Clearing and Settlement

Clearing and settlement constitute a substantial part of the total cost of trading across borders in Europe, especially for retail investors. As such, they are a disincentive for pan-European investment and have been a major obstacle to European market integration. Our concern in this report is with best execution, and here too the high post-trade costs involved in cross-border transactions have, historically, been an important obstruction. This section reviews the technological state of European clearing and settlement systems and discusses the current initiatives and how they are making universal best execution possible throughout Europe.

There are many different contributors to the high cost of cross-border clearing and settlement in the E.U. Many of these, especially those to do with legal and tax issues, are outside the scope of this study and are treated both in the Giovanni Group Report (Giovanni, 2002) and The Comments of the European Central Securities Depositories

Association (ECSDA, 2002). Technical issues are also central, however, and it is important to discuss them here, for as we have argued in earlier sections, the problems facing market consolidation, price transparency and ultimately net price best execution are not exclusively - or even principally - associated with the front office.

7.1 The Need for Integration/Consolidation

Currently clearing and settlement in Europe is very fragmented. There are 19 Central Securities Depositories (CSDs) and three International Central Securities Depositories (ICSDs), Euroclear, Clearstream and SegaInterSettle (SIS). On the clearing side, there are three main Central Counter Parties (CCPs), the London Clearing House (LCH), Clearnet and Eurex Clearing. There is a variety of linkages both between clearing and settlement systems, between CSDs and between ICSDs and CSDs. Figure 7.1, taken from Kramer (2002) is a view of pan-European clearing and settlement as things stood in the second quarter 2002.

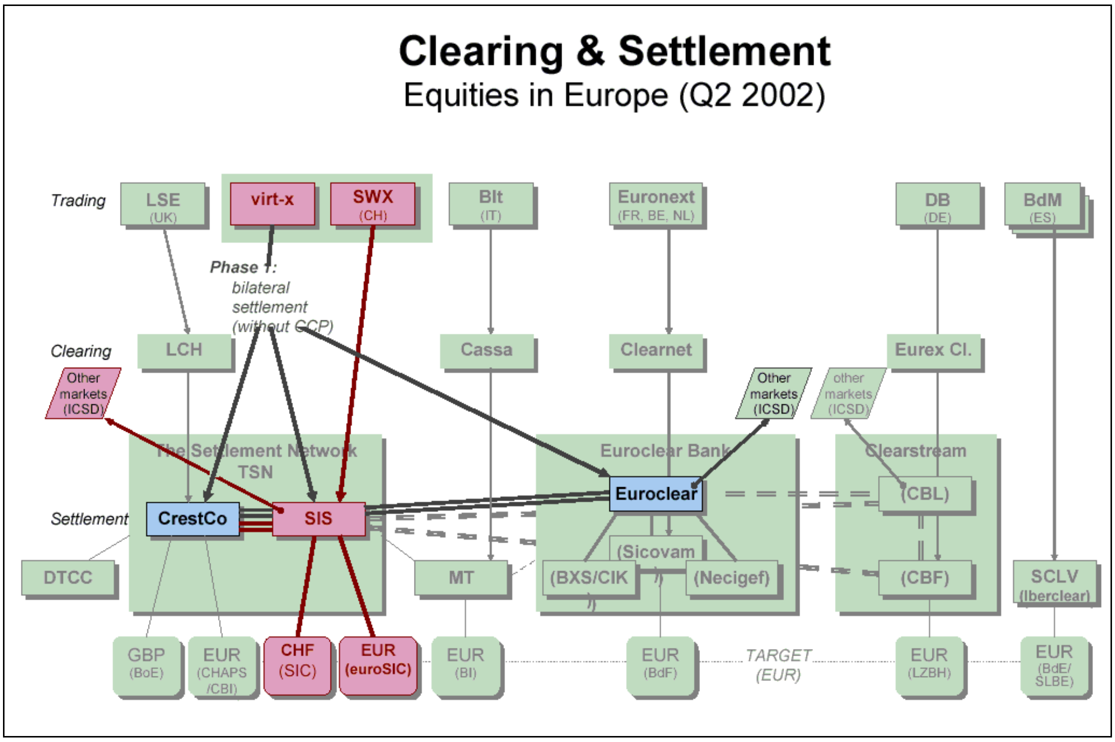

Reproduced from Kramer (2002)

Fig 7.1

There was fragmentation in the USA - although not on this scale - until in 1977 the SEC set up the monopolistic National Securities Clearing Corporation (NSCC) for clearing and the Depository Trust Corporation (DTC) for settlement. These institutions, which merged in 1999 into the Depository Trust and Clearing Corporation (DTCC), have helped to lower clearing and settlement costs in the US.

The inefficiency in the European equities world, caused by the fragmented nature of clearing and settlement, is illustrated by the schematic representation of clearing and settlement interactions required in a typical domestic trade and in a typical cross-border trade shown in Figs 7.2 and 7.3 which are taken from Giovanni (2002).

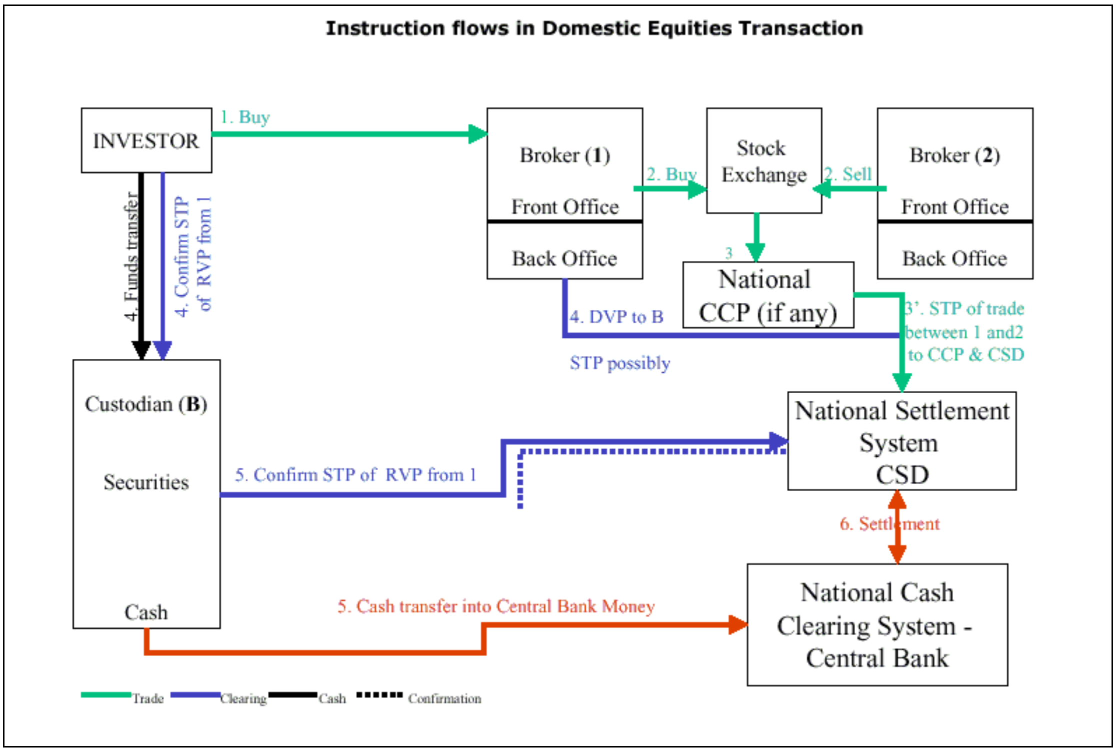

Reproduced from Giovanni (2002) Fig 7.2

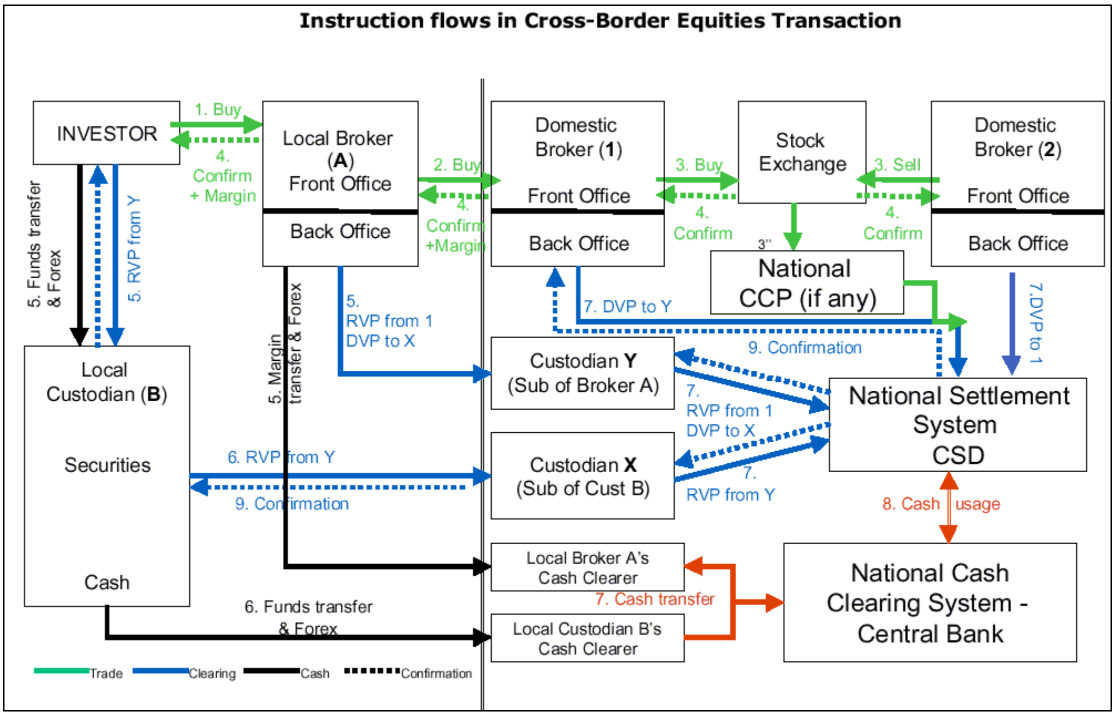

Reproduced from Giovanni (2002) Fig 7.3

Naturally the additional complexity required in cross-border trading creates additional costs and risks for the investor. These arise from the number of intermediaries requiring compensation, the lack of "netting" facilities in some cases and the inefficient and multiple use of collateral. There may also be substantial delays and risks moving securities and cash between systems many of which are manually operated. Estimates vary as to the size of this additional cost. In their study Deutsche Borse/ Clearstream (Clearstream International and Deutsche Borse, 2002) report that their analysis indicates that cross-border wholesale trades cost 30% more and retail trades 150% more than equivalent domestic trades. More striking is the disparity between operating expenditure per transaction incurred by the DTCC and ICSDs. The figure quoted by Levin from an LSE study is more than 5,000 per cent higher for Euroclear over DTCC, and 4,938 per cent higher for Clearstream (Levin, 2001, p. 2).

Clearly there are economies to be made, but, worse from our point of view, the existence of such extreme disparities does tend to undermine the focus on pan-European net price best execution. In interviews, some brokers have told us that at present they consider only the national exchange, for even if there were a better price visible on an alternative exchange they would merely assume that after transferring between depositories the price paid by the investor would be higher, regardless of whether or not this is the case. They base their assumption on the added costs and risk of settlement exceptions involved in such transfers. Whilst for transfers between those CSDs that are still unlinked, this might be a reasonable assumption, for CSDs that are already linked this should not be so.

For the same reason, hedge funds exploiting arbitrage opportunities between stocks quoted on different execution venues tend to unwind the arbitrages by looking for an anomaly in the opposite direction rather than simply resolving the arbitrage by using transfer to a single CSD, because this might currently incur costs prohibitive given the size of the profit opportunity. One consequence of this is that arbitrage is not able to perform its normal role in benefiting efficient price discovery.

7.2 Four Models of Integration/Consolidation

The diversity of IT platforms and interfaces used in clearing and settlement is identified by the Giovanni Group in their report as one of the key barriers to reducing the fragmentation of the European equities market (Giovanni, 2002). Some sort of integration or consolidation would certainly benefit the markets by improving efficiency and reducing the cost of participation for smaller investors in particular. Amongst other benign consequences, we believe this would increase liquidity. Integration is not easily to be achieved, however, and there are competing models to consider.

7.2.1 Single Monopoly Clearing and Settlement House

This option is being promoted by Don Cruickshank, Chairman of the London Stock Exchange (Cruickshank, 2002). The idea is to create a single European depository and clearing house along the lines of the American DTCC. Many however don't believe that this is necessary to produce the cost savings, and others point out that the LSE would benefit from this proposal as the breaking up of the "vertical silos" would be against the interests of the LSE's competitors at Deutsche Borse and Euronext. Political and legal considerations in Europe would clearly make this difficult to achieve on any reasonable time-scale. The current situation, which gives many CCPs an effective monopoly over the exchanges they serve, has led to at least one of them dramatically increasing its fees despite the recent difficult securities market conditions. Very careful regulatory action would be required to avoid this being an even greater problem with a single monopoly.

Perhaps the most crucial objection to this model is the substantial costs to some market participants of adapting their systems and operational interfaces to the platform of the merged utility. Not only are these costs high but, since they fall only on participants who are not already interfaced to the CSD in question, whereas the benefits are experienced by all participants, they are discriminatory. Problems of this sort were felt acutely in Germany when Auslandskassenverein was transferred to the Creation platform of Clearstream: one market participant, in particular, calculated the payback time for adapting its internal systems at more than twelve years.

7.2.2 ICSD with Hub and Spoke Architecture

At one time this was the favoured option of Euroclear who already have a lot of the links to national CSDs in place. The idea is that national CSDs handle domestic settlement and pass cross-border trades on to a single central ICSD. Some doubt that without competition at the ICSD level cost savings would be all that attractive to retail investors, and it is certainly not obvious that the best system would be the winner in an environment where national interest backed by money could heavily influence the outcome. More fundamentally, the fact that the ICSD "hub" is in competition with the CSD "spokes" makes this model unattractive to the participating CSDs which is one of the issues that has, historically, proved problematic in the adoption of "hub and spoke" architecture in the FX world.

7.2.3 Multiple CSDs with Bilateral Linkages

Linkages between CSDs may be used effectively to "domesticate" cross-border trades. Already Crest has linked with the Swiss depository SIS to offer sale and receipt of UK, Irish, Swiss and other major European stocks listed on virt-x against payment in sterling, euros or US dollars. The ECSDA has promoted this general approach and has drafted proposals for standardisation of the communications between CSDs. These are based on the globally accepted ISO 15022 standard defined by the International Organization for Standardization (ISO) (ECSDA, 2002, p11). Nevertheless, there is some opposition to this approach. Levin, for instance, argues from the suggestion that n * (n - 1) bilateral linkages would be required, that it would be neither practical nor cost-effective to link the CSDs in this way (Levin, 2001). We do not find this argument compelling. As always, we make a distinction between logical and physical links. If different logical links were required for each pairing then the proposal would be expensive and impractical. If, however, standardised logical links are used, as proposed by the ECSDA among others, the incremental cost of physical links need not be high.

More problematic to any approach relying on bilateral linkages is the question of what to do with securities of a third country that are involved in a link between CSDs of two other nationalities. Such transactions may continue to experience greater complexity and cost.

7.2.4 Communications Hub with CSD Spokes

n * (n – 1) linkages would *not* be needed if the network made use of "peering" centres as is typical with internet linkages. Tier 1 Nexgen Peering centres have proved effective in addressing the major problems associated with bilateral or private peering, including physical complexity , telecommunications costs, bandwidth, and speed of provisioning of new links. Fig 7.4 based on Linton (2002) illustrates this development.

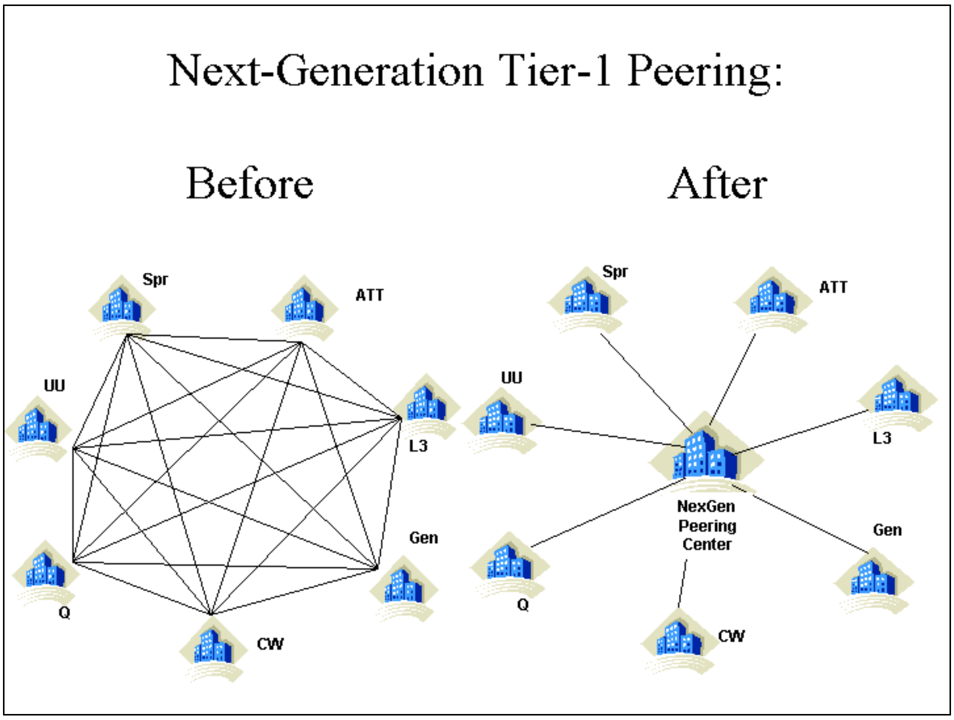

Adapted from Linton (2002)

Fig 7.4

This model confers the benefits of the "hub-and-spoke-model" without the intermediation of an ICSD which was the fundamental weakness of that model. We recognise that standardisation of logical links including messaging is critical to the success of such a model.

One proposal along these lines is the current Central Securities Settlement Infrastructure (CSSI) initiative between SIS, Crest, Monte Titoli, IberClear and the DTCC (see Figure 7.5, reproduced with permission from a private document). This has a central hub which is not a CSD or CCP, so is not in competition with the CSD "spokes".

The model is designed to maximise "settlement liquidity" by settling each equity in the CSD where costs are lowest - normally the domestic CSD. To facilitate this, the plan is for participating CSDs to implement minimum "service level standards" which go beyond ISO 15022 and simple messaging standards. The idea, put simply, is that if CSD1 sends a message to CSD2, all things that CSD1 expects to happen at CSD2 do happen, though not necessarily in the way that they would at CSD1 itself. The end result of all flows of information, cash, equity and collateral is the same.

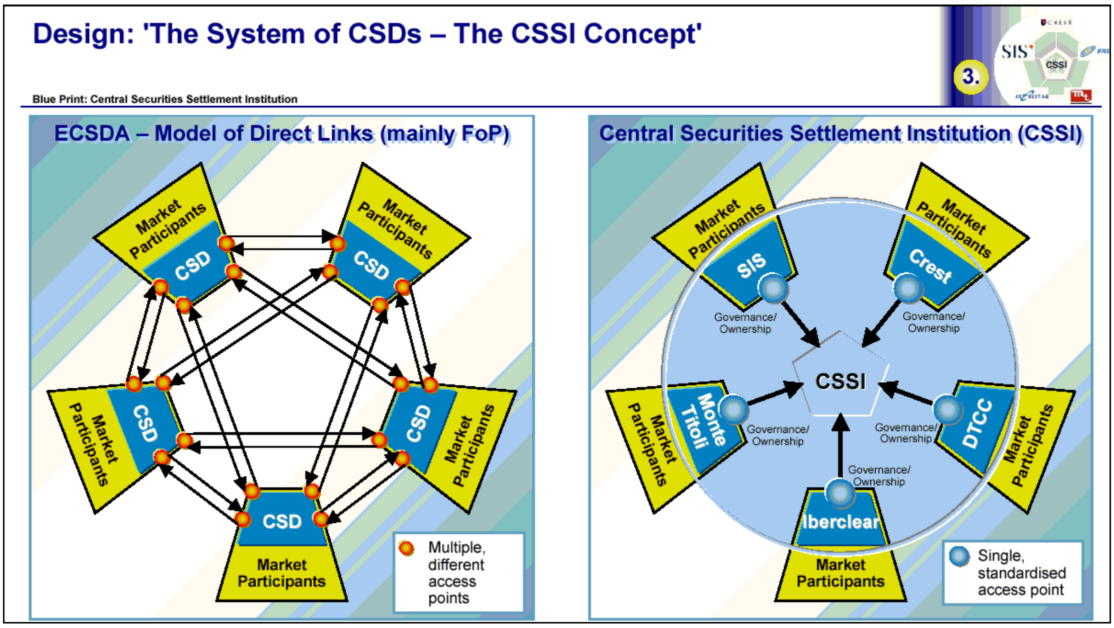

Reproduced from a CSSI presentation - © SIS 2002 Fig 7.5

The interviews that we have undertaken have shown a wide range of middleware and private network vendors are ready to provide all the software and physical connectivity necessary to implement efficient networks removing the need for bilateral linkages.

7.3 Costs

From the CSD's point of view, we believe that the additional cost to them of providing the necessary linkage would be compensated by revenues flowing from increased volumes of cross-border trading that would be expected to follow, even allowing for anticipated reductions in clearing and settlement charges.

Turning to the customer: we address the issue of increased customer cost associated with cross-border traffic from a technical point of view. We note that, in internet service provision, although smaller networks typically pay larger networks for connectivity, "many large networks themselves exchange their traffic without charge under a peering, sender-keeps-all basis" (TeleGeography, 2000). This has been commonplace both amongst the largest Tier-1 providers (at peering centres) and local Internet Service Providers (ISPs) at local exchange points, and has proved simple and cost-effective for customers. "Sender keeps all" is unlikely to be an appropriate charging structure for this model of CSD integration, but the use of standardised cost/benefit analysis is now usual in the ISP "settlement-free peering" versus "paid transit" decision process (Norton, 2001), and similar techniques should be applicable for linked clearing and settlement solutions.

7.4 Straight Through Processing

Horizontal integration, along the lines proposed here, need not compromise efforts towards Straight Through Processing (STP)[11] as some respondents to the Charteris survey

[11] We define Straight Through Processing (in its ultimate form) as fully automated processing of an order entered by a customer through broker, exchange, clearing house, depository and payment systems. STP is

suggested (Charteris, 2001, p. 32). Properly designed, STP should work with fully described data and objects, which should entail no trade-off between STP and integration. It has long been held as a good software design principle to minimise the "coupling" between separate modules. Stevens et al (1979, p29) define "coupling" as "the measure of strength of association established by a connection from one module to another. Strong coupling complicates a system since a module is harder to understand, change, or correct by itself if it is highly interrelated with other modules. Complexity can be reduced by designing systems with the weakest possible coupling between modules".

Major initiatives in financial and technological arenas are underway to link disparate systems by way of XML protocols with internationally agreed schemas including standardised data vocabularies (see next Section), indicating that loose coupling is now widely accepted in the financial industry. This is the complete opposite of vertical integration with proprietary formats - a separate STP for each vertical "silo" is a consolidator's nightmare which would lead to even greater costs and unnecessary duplication of effort.

7.5 Different Settlement Cycles and Other Local Anomalies

Although it has been necessary to consider clearing and settlement issues in more detail than we expected when we embarked on this report, we should remind the reader that our primary focus is best execution and technological issues relating to it. Nevertheless we shouldn't leave the issue of clearing and settlement without referring to the obstacles to pan-European integration posed by local anomalies, an example of which, mentioned by both the Giovanni Group and the ECSDA, is the difference, currently existing from jurisdiction to jurisdiction, between settlement cycles. It would be outside the scope of this report to consider this or other local anomalies in detail, but we draw the reader's attention to the recent soft launch of the Continuous Linked Settlement system (CLS, 2002). It appears that with CLS the FX world may have finally solved its complex, multipartite, 24 hour, global netting arrangements, which gives us optimism that the global equities world can follow suit[12].

7.6 The Role of Regulation in Mandating Linkages

Without regulatory intervention it is extremely unlikely that *all* the necessary linkages to remove the barriers to universal pan-European cross-border trading will come about. The Charteris report (Charteris, 2001) argues that the major investment banks which dominate the European equities business prefer the status quo. In particular, they observe that banks' current systems handle the existing complexity which provides an effective barrier to entry, and that costs of the inefficiency are simply passed on to the customer with a mark-up.

We accept this analysis and believe that linkages between clearing and settlement systems need to be mandated. We have argued that the costs to the CSDs themselves need not be severe, and we believe the benefits to the overall market would be enormous in terms of

---

often understood in a more restricted sense limited to the automated processing of fills by back office systems, which is the way we used the term ourselves in Fig 5.1.

[12] For a currency to be CLS eligible that currency has to operate an RTGS - a real-time gross settlement system - meaning that there are no longer settlement cycles. See De Feo (2002) for details.

market integration, the promotion of cross-border trading, increased liquidity, market efficiency and greatly reduced costs to retail and institutional investors.

Integration of the European securities and settlement infrastructure is being promoted by the collaboration between the European System of Central Banks and the Committee of European Securities Regulators (ECB/CESR, 2002) formed specifically for this purpose, and is also being addressed in practice as we have seen by initiatives between some providers. These initiatives should mean that the historically crippling post-trade costs of cross-border trading, which have been so dominant in deciding where the best net price is available, will become much less so in the near future. This makes our attention to best net price execution very much more apposite. Indeed, it is the authors' opinion that development of the regulatory framework to include net price based best execution needs to start now to keep pace with the rapid developments in the linkage of clearing and settlement systems.

7.7 Conclusions

7.7.1 The fragmented nature of clearing and settlement in Europe has been a major obstacle - perhaps *the* major obstacle historically - to market integration and best execution.

7.7.2 Already CSD linkages have "domesticated" some cross-border trading.

7.7.3 After examining the costs involved in different models of integration/consolidation we have shown that at least some of these models offer a viable pan-European architecture for linking clearing and settlement systems.

7.7.4 Initiatives already exist between some CSDs to implement such a model.

7.7.5 We believe that regulatory intervention will be necessary to *complete* these linkages across all European venues, because the status quo favours those large firms whose systems already cope with the complexity, and who are able to pass on the costs of inefficiencies with a mark-up, ultimately to investors.

7.7.6 The costs incurred by the industry in implementing these linkages should be more than compensated by the increased volume of cross-border trading, even allowing for consequent reduced clearing and settlement charges per transaction.

7.7.7 Rapid development in linkage of clearing and settlement systems means that development of the regulatory framework to include pan-European net price based best execution needs to start now.

8 Exchange Interfaces

Best execution requires reliable price consolidation from multiple execution venues. As we have seen, this has historically been hampered in Europe by the multiplicity of exchange interfaces. In this section we examine the need for standardisation of these interfaces, analyse the form that it might take and the issues that will need to be addressed, and discuss the current initiatives in this area.

8.1 The Need for Standardisation

As already set out in section 1 the financial industry is subject to substantial costs and barriers to competition caused by each exchange/execution venue having its own interface[13] for transmission of pre- and post-trade data. Consequently there are relatively few international consolidators and a multiplicity of domestic vendors offering very partial coverage. The costs are so high because, even having interfaced, a consolidator must retain significant resources to cope with regular changes to the feeds, which leads to substantial and ongoing duplication of interfacing efforts across the financial industry.

To achieve effective pan-European market integration and universal availability of low-cost net price benchmarking it would be advisable to reduce the resource requirements and lower the barriers to entry. Perhaps surprisingly, our interviews have shown that at least some major consolidators are not opposed to an international standardisation of formats as they believe this will free their resources to compete on value-added services. Such standardisation will have a number of effects, reducing the costs to the financial industry which are borne directly or indirectly by investors, facilitating faster development of financial systems and removing barriers to market integration.

Similar arguments can be applied to the standardisation of the basic order routing interfaces, and in the case of equities this is happening. The Financial Information eXchange protocol (FIX) has already achieved substantial take-up in both the US and Europe as a mechanism for passing all manner of equity orders and fills. Furthermore, ISO's working group 10 (ISO 15022 XML WG10) is in the process of unifying the FIX and SWIFT (Society for Worldwide Interbank Financial Telecommunications) protocols into one XML (known as ISO 15022 XML), which is also intended to cover the Financial Products Markup Language (FpML), in addition to the existing ISO 15022 protocol.

> *Specifically, FIX Protocol is contributing its expertise in the pre-trade/trade execution domain, and SWIFT will provide post-trade domain expertise to the ISO 15022 XML efforts. SWIFT is also acting as the Registration Authority. In this capacity, it maintains the Data Field Dictionary/Catalogue of Messages, supports the Registration Management Group, and monitors the Registration Authority.*
>
> *The aim is to migrate the securities industry to a standardized use of XML to ensure interoperability across the financial industry (*XMLonWallStreet, 2002).

However, whilst many execution venues have taken up at least some version of FIX for order routing, to date no internationally agreed standard for real-time transmission of pre-trade quotes and post-trade data has been available. As we reveal in Section 8.4, this is about to change.

[13] As stated in Section 1.2.1, we understand "interface" broadly to encompass logical, protocol, messaging, API, vocabulary and data field definition issues.

8.2 XML Standards

XML is achieving widespread acceptance as a preferred means of linking heterogeneous systems worldwide. The World Wide Web Consortium (W3C) was created in October 1994 to lead the World Wide Web to its full potential by developing common protocols that promote its evolution and ensure its interoperability. W3C describes XML in its Activity Statement (World Wide Web Consortium, 2002) as a "simple, very flexible text format derived from Standard Generalized Markup Language (SGML, ISO 8879). Originally designed to meet the challenges of large-scale electronic publishing, XML is also playing an increasingly important role in the exchange of a wide variety of data on the Web".

It goes on to list perceived benefits, which include:

- Saves business money by enabling the use of inexpensive off-the-shelf tools to process data
- Saves training and development costs by having a single format for a wide range of uses
- Increases reliability, because user agents can automate more processing of documents they receive
- Encourages industries to define platform-independent protocols for the exchange of data, including electronic commerce
- Allows people to display information the way they want it, under style sheet control
- Enables long-term reuse of data, with no lock-in to proprietary tools or undocumented formats

All of these are relevant to concerns in the financial data arena.

Typically a working group is formed to establish a new XML protocol (e.g. ISO 15022 XML) for a particular arena. The definition of the XML protocol is embodied nowadays in an XML schema. The schema sets out the shared vocabulary and also defines the rules governing the possible structure, content and semantics of all XML documents using that XML protocol. Appendix B illustrates this with an example XML document and its underlying XML schema for a theoretical "Purchase Order XML Protocol".

XML parsers (unlike conventional HTML parsers used to display web pages) are required to reject as "not valid" any XML document which violates the schema for the protocol in any way. This significantly reduces the probability of errors propagating between systems linked by XML protocols.

Namespaces allow the use of multiple schemas within a single XML document or feed. Extensibility is therefore a key feature of XML.

8.3 Addressing XML Issues

XML is typically verbose by design and can be slow to parse. This is often seen as a barrier to the use of XML protocols for high-frequency, high-volume inter-process communications. Indeed as an extreme example, the authors were once supplied with a prototype "real-time" market data XML feed by a major institution which transmitted over 150 characters per quote (1200 bits).

This perceived barrier will be lowered by progress on both hardware and software fronts. On the hardware side, a significant development is the availability of network devices capable of intelligent XML handling, including compression and routing across communications links. Devices of this kind are expected to become widespread in the future, which will greatly improve the efficiency and speed of XML communications across geographically diverse (wide area) networks and systems. Forrester Research - a firm that identifies and analyses emerging trends in technology and their impact on business - forecasts that Web services will significantly increase the amount of XML traffic on networks, and have recently published a brief entitled "Hardware Help for the Looming XML Blitz" on this topic (Forrester, 2002).

On the software side, compression is also addressing verbosity. There is active research into semantic compression algorithms which can take advantage of the specific syntactic and semantic properties of text strings in the XML documents. Algorithms such as the Burrows-Wheeler Transform are being modified to improve performance on text and XML documents (Chapin and Tate, 1998). Software tools, designed for XML compression, such as Xmill are also widely available now. Results published by the developers of Xmill show that when compressing XML it considerably outperforms conventional (non-XML specific) compressors both in terms of compression ratios and speed of compression (Liefke, 1999).

Speed of transmission and parsing is also being successfully addressed. Some commercial vendors already offer XML-based middleware which they claim is suitable for inter-process communication of real-time market data at least on local networks. It is important to note that, at present, local area networks usually have 50 to 500 times the available bandwidth of wide area networking (including exchange connectivity). This means that speed of processing rather than verbosity is key for XML middleware, whereas terse design and compression ratio are critical where XML is used for exchange interfaces.

Our own company currently uses proprietary compressed XML middleware to facilitate distributed processing. This includes an "adapter-driven format" (adaptable format stored in a database table) for compressed high-frequency market data, which has been benchmarked at up to 60,000 quotes per second between two mid-range PCs.

8.4 The Market Data Definition Language

There is a published XML standard for market data known as MDDL (Market Data Definition Language), which is specified under the auspices of the Financial Information Services Division of the Software and Information Industry Association (FISD) and intended as an international standard. The official website describes MDDL as *an XML-based interchange format and common data dictionary on the fields needed to describe 1) financial instruments, 2) corporate events affecting value and tradability and 3) market-related, economic and industrial indicators. The principal function of MDDL is to allow entities to exchange market data by standardizing formats and definitions. MDDL provides a common format for market data so that it can be efficiently passed from one processing system to another and provides a common understanding of market data content by standardizing terminology and by normalizing the relationships of various data elements to one another.* (MDDL, 2002)

To date MDDL has been concentrated on snapshot, reference and end of day data. However, our research shows there are plans to address real-time market data in 2003. The view of the FISD, expressed in private communications, is that market data is suitable for specific compression that generic compression schemes cannot address. Highly relevant here are Liefke's experiments showing outperformance, in both space and time, of Xmill using the option for user compression (specific compression for a particular XML protocol) over Xmill with default compression (Liefke, 1999). The FISD's working assumption, which has been used to analyse the XML requirements for real-time data feeds, is that it should be possible to deliver 50,000 transactions per second, from a universe of 1,000,000 instruments, over a pipe of 1.5 Megabits per second, i.e. 30 bits per transaction.

The authors' own experience shows this data rate to be a reasonable objective, although we note that in addition to bandwidth considerations, conventional XML parsers could be overwhelmed by the requirements of processing high-frequency bursts of real-time market data. We also note that current non-XML data feeds sometimes handle updates of real-time order books by passing the volumes on all of the N nearest bids and offers on each update. This can greatly increase the size of updates and hence bandwidth requirements. It is not unknown for these issues to be addressed, in the case of existing non-XML data feeds, by throttling[14] the updates, with consequences for the completeness and timeliness of the market data streams. Complications of this sort are best avoided in the construction of a reliable net price benchmark, and fortunately they are relatively easy to address in XML and non-XML data feeds alike, for instance, by passing only the changes in data, with periodic refreshes of the unchanged data.

The FISD point out that the XML reporting format needs to be sufficiently flexible for a provider to state exactly what a data field is. As a simple example they note that if the data feed field is "change" then the XML should describe exactly what that change is – is it the change since the "open" or the change since the previous "close"? They further argue that providers should not be limited as to *what* they can provide, but only limited as to *how* they provide it - namely, in a standard cross-industry format, or, more specifically, using the exact same vocabulary. They class MDDL as an effort to describe a common vocabulary at an international level *and* to have a common XML to deliver data consistent with those definitions.

[14] By "throttling" we understand the limiting of bandwidth over a telecommunications link. To cope with bursts of activity, which would otherwise violate the bandwidth requirements, exchange data may be "throttled" either by queuing which introduces time delays or by sampling which removes data from the feed.

8.5 Mandating Minimum Standards

We agree that such an effort should result in common logical linkages and go a long way towards resolving the problems posed by the current "tower of Babel" if adoption is widespread. However we believe that in pursuit of simple calculation of a pan-European price benchmark a further step needs to be taken. This is to mandate the timely supply of a *minimum common set of data fields* to a pan-European standard for pre- and post-trade data by execution venues, via an agreed XML standard. Execution venues could still supply additional data feeds or schemas where they feel this adds value but the additional data feeds should not supply data in advance of that supplied in the "mandated minimum feed".

A common minimum set of shared *meanings* is key, and the distinction between this and the corresponding shared *vocabulary* (or set of words) is crucial:

> *If an application concentrates on the word rather than the meaning, there is a risk that separate meanings … may end up being used interchangeably - with disastrous consequences* (LMS, 2002).

The "change" data field example in 8.4 above is a case in point.

Consolidators rather than execution venues suffer the bulk of the current interfacing and maintenance costs and they pass these on to consumers in the financial industry. As a result, there is little impetus to resolve the current "babble" of exchange data and its consequent costs to investors, barriers to competition, and massive duplication of efforts. Sensitively applied, regulatory action to establish a mandated minimum feed should lower the costs to the industry as a whole, and to its consumers, whilst encouraging market integration. Such changes are important to increase the number of vendors able to offer pan-European best execution modules.

8.6 Conclusions

8.6.1 The current multiplicity of exchange interfaces, for both pre- and post-trade data, results in substantial and ongoing duplication of effort in the financial industry, inflation of its cost base and a barrier to competition.

8.6.2 These increased costs are ultimately passed on to the customer, so the issue of exchange interface standardisation is important to customer protection.

8.6.3 XML standards are emerging, incorporating standardised vocabularies and data field definitions, and supported by hardware and software technologies, which are capable of simple loose coupling of *real-time* financial systems.

8.6.4 Because of the structure of the industry there exists no incentive for execution venues to agree to, and implement, one of these standards, so regulatory action will be necessary to require them to provide a minimum common set of data fields ("vocabulary") in a standard cross-industry format, and that they provide this in timely fashion.

8.6.5 An exchange should be at liberty to provide any additional data feeds that it might wish, but these should not be provided in advance of the "mandated minimum feed".

8.6.6 These regulatory moves should reduce costs to the industry and, ultimately, to investors. They would also promote market integration and facilitate competition amongst vendors to provide pan-European best execution modules.

9 Main Conclusions and Recommendations

9.1 Main Conclusions

9.1.1 Technology already exists capable of performing pan-European consolidation of quotes, order books and trades.

9.1.2 Technology already exists to identify in real-time the best available price net of all explicit costs (including commission, exchange transaction fees, taxes, clearing, settlement) and allowing for currency conversion where necessary.

9.1.3 Technology already exists to route orders to the European execution venue offering the best net price.

9.1.4 A service combining best net price calculation on multiple markets with routing to the market offering the best price could be offered very inexpensively on an Application Service Provider model.

9.1.5 The US experience has shown that a benchmark, though imperfect and easy to beat, can be of great value in providing investor protection, and is a prerequisite for the emergence of price improvement services.

9.1.6 The pan-European best net price is a suitable benchmark for European best execution of market orders and marketable limit orders of a size related to the total displayed interest in the market.

9.1.7 The fragmented nature of clearing and settlement in Europe, and the variety of interfaces (including data formats and protocols) used by the various execution venues, have been the principal obstacles to market integration and best execution.

9.2 Main Recommendations

9.2.1 The pan-European best net price benchmark should be mandated. Our suggestion is that brokers be required to direct customers' market orders (of a size not exceeding the total displayed interest) to the execution venue offering the VWABNBO price to that size, or, alternatively, offer execution at a price at least as good. Sufficient contemporaneous data should be recorded to allow monitoring of compliance with this obligation.

9.2.2 The completion of linkages between the different clearing and settlement systems in Europe should be mandated. European initiatives are in progress, but to achieve maximum benefits from 9.2.1 (along with the other advantages of integration) we see a need for regulatory pressure to ensure early implementation

9.2.3 A standardised minimum feed should be mandated to cover all European execution venues to provide pre- and post-trade data in timely fashion. Our suggestion is that an XML schema, including a minimum common vocabulary, be defined and its implementation by all European execution venues be required by the regulators.

9.3 Impact of Suggested Outline Regulatory Package

A regulatory package along the lines suggested in 9.2 would promote the use of available technology and facilitate competition between execution venues, consolidators, technology vendors and CSDs. We would expect this to benefit market integration, price discovery and customer protection, and to drive down the overall cost of trading. It is possible that sensible harmonisation of trading tax regimes across the European market could also emerge in this way without being imposed, as a consequence of national self-interest.

If such a package were implemented we would expect the overall cost savings to the industry and consumers to exceed by far the implementation cost. Furthermore, increased investor confidence, together with the lowered costs, should promote trading activity, giving rise to higher volumes and improved market liquidity, which would be in everyone’s interests.

The authors are strongly opposed to the imposition of unnecessary regulation, but the nature of the market and vested interests involved make it extremely unlikely that the changes discussed here could be completed without regulatory backing.

Moves towards standardising exchange interfaces and linking clearing and settlement systems are of value in their own right, but they are of particular importance in removing obstacles to the adoption of a best net price standard for best execution, which is our principal focus. That is why we propose that our outline regulatory recommendations be implemented as a package. Pan-European trading needs a “cornerstone” to stand alongside the US National Market System, to provide clear and enforceable customer protection to European equities investors, and to drive the integration of the European securities market so that it can more effectively compete for investment flows in an increasingly global economy.

Bibliography

Alba, J. N. (2002) "Transaction Costs Analysis: How to Achieve Best Execution", Best Execution: Executing Transactions in Securities Markets on Behalf of Investors, European Asset Management Association, Heronsgate, http://www.eama.org/publications/bestexecution.pdf, p 6-17.

AIMR (2001), "Trade Management Guidelines", Association for Investment Management and Research, http://www.aimr.com/pdf/standards/proposed_tmg.pdf.

Association of Private Client Investment Managers and Stockbrokers - European Association of Securities Dealers, et al (2002), "Innovation, Competition, Diversity, Choice: A European Capital Market for the 21st Century", Position Paper, http://www.isda.org/speeches/pdf/Innovation_Competition_Diversity_Choice.pdf.

Bacidore, J., Ross, K. and Sofiano, G. (1999), "Quantifying Best Execution at the New York Stock Exchange: Market Orders", NYSE Working Paper 99-05.

Battalio, R., Hatch, B. and Jennings, R. (2000), "Dimensions of Best Execution for Market Orders: Assessing Differences between the NYSE and the Nasdaq Third Market", http://www.bus.indiana.edu/Finance/workingpapers/NYSECON5.pdf.

Battalio, R., Jennings, R. and Selway, J. (2000), "The Relationship between Market-making Revenue, Payment for Order Flow, and Trading Costs for Market Orders", http://www.bus.indiana.edu/Finance/workingpapers/knight13.pdf.

Board, J. and Wells, S. (1999), "Liquidity and Best Execution in the UK: A Comparison between SETS and Tradepoint", http://fmg.lse.ac.uk/events/conferences/pastcon/exchanges/papers/board&wells.pdf.

Chapin, B. and Tate, S. F. (1998), “Higher Compression from the Burrows-Wheeler Transform by modified sorting”, Dept of Computer Science University of North Texas, http://www.cs.unt.edu/~srt/papers/bwtsort.pdf

Charteris (2001), "European Financial Markets 2001: Revolution or Evolution", http://www.charteris.com/Reports/EFM/ExchangesDownload.asp.

Clearstream International and Deutsche Borse (2002), "Cross-Border Equity Trading, Clearing and Settlement in Europe", White Paper, http://deutsche-boerse.com/INTERNET/EXCHANGE/zpd.nsf/Web+Publikationen+E/AKLS-58TMV6/$FILE/White-Paper_online_e.pdf?OpenElement.

CLS (2002), Overview Presentation of Continuous Linked Settlement, http://www.cls-group.com/news/resource.cfm?id=AB4F2CEB-15BD-405B-9727DF2E7BEF0A3D

Committee of European Securities Regulators (2002), "A European Regime of Investor Protection: The Professional and the Counterparty Regimes", Ref: CESR/02-098b, http://www.europefesco.org/DOCUMENTS/STANDARDS/02-098b.pdf

Crest (2002), "Enhancing CREST: September 2002 - Enhanced Delivery versus Payment and Free of Payment link to SIS" Version 2, White Paper, http://www.crestco.co.uk/publications/white_books/enhancing_sis_0902.pdf

Cruickshank, D. (2002), "Clearing and Settling European Securities: Where Competition Works (and Where it Doesn't)", Speech to European Harmonisation of Cross-Border

Settlement Conference 21/22 March 2002, http://www.londonstockexchange.com/newsroom/speeches/speech15.asp.

De Feo, J. (2002), “Continuous Linked Settlement - Shaping the Future of Settlement”, IFI Magazine, Spring 2002, http://www.cls-group.com/news/article.cfm?objectid=C3619C10-A06F-4875-B029F648176FB7E4.

Deutsche Borse (2001), Implementation Regulations of the Frankfurt Stock Exchange Concerning Technical Equipment for the Electronic Trading System, (Deutsche Borse Group July 30, 2001), http://deutsche-boerse.com/INTERNET/EXCHANGE/zpd.nsf/Web+Publikationen+E/HAMN-52CEW3/$FILE/FWB03-01-07-30e.pdf?OpenElement

Domowitz, I., Glen, J. and Madhavan A. (2001), “Global Equity Trading Costs”, http://www.itginc.com/research/whitepapers/domowitz/globaleqcost.pdf

ECB/CESR (2002), Joint work of the European System of Central Banks and the Committee of European Securities Regulators in the Field of Clearing and Settlement, Call for Contributions, Ref.:CESR/02-005b.

ECSDA (2002) , "Comments of the European Central Securities Depositories Association regarding the Communication from the Commission to the Council and the European Parliament 'Clearing and Settlement in the European Union - Main Policy Issues and Future Challenges'", http://www.ecsda.com/C125666A0027EF89/Pages/FramesIE?opendocument

Eurex (2002), "Guidelines for Exchange Participants Eurex 2002", http://www.eurexchange.com/eic/download/NCMguideline2002en.pdf.

Euronext (2002a), "Internalisation", Euronext internalisation proposal, http://www.bondmarkets.com/europe/euronext_proposal.pdf.

Euronext (2002b), Revision of Investment Services Directive, Euronext’s Proposals to the ISD in relation to the EU’s cash equity markets, http://www.bourse-de-paris.fr/en/news7/pdf/Position_paper_internalisation_24052002.pdf.

FSA (2002), "Financial Services Authority Consultation Paper on Best Execution", Consultation Document 154, http://www.fsa.gov.uk/pubs/cp/cp154.pdf.

Forrester (2002), Forrester Brief, "Hardware Help for the Looming XML Blitz," March 22, 2002.

Giovanni (2002), Report of the Giovanni Group, "Cross-border clearing and settlement arrangements in the European Union", http://europa.eu.int/comm/economy_finance/publications/economic_papers/2002/ecp163en.pdf.

Harris, L. E. (1996), "The Economics of Best Execution", http://www-rcf.usc.edu/~lharris/ACROBAT/BESTEXEC.PDF.

Kramer, G. (2002), "The European Stock Exchange and Settlement Landscape", Slides from presentation, Madrid 25 April 2002, http://www.swissplus.ch/madrid_020425_kramer_telekurs_e_pp.pdf.

Levin, M. (2001), "Clearing and Settlement in the European Union", CEPS Commentary, http://www.ceps.be/Commentaty/July01/csd.php.

Levitt, A. (1999), "Best Execution: Promise of Integrity, Guardian of Competition", Speech by the SEC Chairman to Securities Industry Association, Boca Raton, Florida, November 4, 1999, http://www.sec.gov/news/speech/speecharchive/1999/spch315.htm.

Levitt, A. (2000), "Best Execution, Price Transparency and Linkages: Protecting the Investor Interest", Washington University Law Quarterly, 78, 513-8, http://law.wustl.edu/WULQ/78-2/levitt.pdf.

Levitt, A. (2001), "The National Market System: A Vision that Endures", http://www.sec.gov/news/speech/spch453.htm.

Liefke, H. (1999), "Xmill – an efficient compressor for XML", http://www.research.att.com/sw/tools/xmill/experiments.html

Linton, J. R. (2002), "New Directions in Peering for Tier 2 and Content Providers", NANOG 23, www.nanog.org/mtg-0202/ppt/linton.ppt

LMS (2002), XML Standards Guide, Semantic Dimensioning, Official Website of London Market Systems, http://www.londonmarketsystems.com.

McCleskey, S. (2001), "Best Execution in Cross-Market Trading Environments", Dissertation for the M.Sc. Financial Regulation and Compliance Management at London Guildhall University.

McCleskey, S. (2002), "Best Execution in a Cross-Market Environment: The Case for a Pan-European Policy", The Journal of International Financial Markets, Volume 4, Issue 1, pp. 1-5.

MDDL (2002), Official Website of the Market Data Definition Language, http://www.mddl.org.

Myners Report (2001), "Institutional Investment in the United Kingdom: A Review", http://www.hm-treasury.gov.uk/mediastore/otherfiles/31.pdf.

Norton, W. B. (2001), "A business case for ISP Peering", http://www.cs.berkeley.edu/~randy/Courses/cs294.s02/Case4Peering.pdf.

O'Hara, M. (2001), "Overview: Market Structure Issues in Market Liquidity", BIS Papers No. 2 - Market liquidity: proceedings of a workshop held at the BIS, http://www.bis.org/publ/bispap02a.pdf.

Schwartz, R. A. and Steil, B. (2001), "Controlling Institutional Trading Costs: 'We Have Met the Enemy, and They Are Us'", Journal of Portfolio Management, Spring 2002, Volume 28, Number 3, Pages 39 - 49.

SEC (2000), "Final Rule: Disclosure of Order Execution and Routing Practices", 17 CFR Part 240, [Release No. 34-43590; File No. S7-16-00], RIN 3235-AH95, http://www.sec.gov/rules/final/34-43590.htm.

SEC (2001a), "Report of the Advisory Committee on Market Information: A Blueprint for Responsible Change", http://www.sec.gov/divisions/marketreg/marketinfo/finalreport.htm

SEC (2001b), Release 34-43963, http://www.sec.gov/litigation/opinions/34-43963.htm.

SEC (2002), "Order with Respect to the Implementation of NASDAQ's SuperMontage Facility", Release No. 46429, File No. SR-NASD-99-53, http://www.sec.gov/rules/sro/34-46429.htm.

Selway, J. P. (2002a), "Five Myths about Listed Trading", Institutional Investor, Spring 2002, http://www.institutionalinvestoronline.com/iiochannel/institutionaltrading/TP_SP_02_Selway.pdf.

Selway, J.P. (2002b), "SuperDF, Alternative Montage, and the S-Word: A Hitchhiker's Guide to Complexity", Wall Street and Technology, http://www.wallstreetandtech.com/story/WST20020617S0001.

SIS and Crest (2001), "TSN - The Settlement Network", The Guide to European Equities 2001, http://www.thesettlementnetwork.com/tsn_euromoney.pdf

Smant, D. J. C. (2002), "Financial Market Micro Structure", Lecture Notes Workshop/Seminar Monetary Economics, A7149 Interest Rates and Stock Markets , http://www.few.eur.nl/few/edu/vak/a7149/stdmat/2bmicromrkt_.pdf.

Stevens, W., Myers, G. and Constantine, L. (1979), "Structured Design", in Classics in Software Engineering, New York, NY, Yourdon Press.

TeleGeography (2000), "Hubs and Spokes: A Telegeography Internet Reader", ISBN 1-886142-23-8.

USCA (2001), United States Court of Appeal 3rd Circuit, No. 00-1586, "Newton v Merrill Lynch", http://www.ca3.uscourts.gov/opinarch/001586.txt.

W3C (2001), World Wide Web Consortium, XML Activity Statement, XML Schema Part 0: Primer W3C Recommendation, 2 May 2001, http://www.w3.org/TR/xmlschema-0)

W3C (2002), World Wide Web Consortium, XML Activity Statement, www.w3c.org/XML/Activity.

XMLonWallStreet (2002), "Financial XML projects ISO 15022 XML", www.xmlonwallstreet.com

traderserve

Appendix A

Abbreviations

ADF	Alternative Display Facility
API	Application Programming Interface
ASP	Application Service Provider
CCP	Central Counter Party
CESR	Committee of European Securities Regulators
CLS	Continuous Linked Settlement
CRM	Customer Relationship Management system
CSD	Central Securities Depository
CSSI	Central Securities Settlement Infrastructure
DTC	Depository Trust Corporation
DTCC	Depository Trust and Clearing Corporation
DVP	Delivery Versus Payment
ECN	Electronic Communications Network
ECSDA	European Central Securities Depositories Association
FISD	Financial Information Services Division (of the Software and Information Industry Association)
FIX	Financial Information eXchange protocol
FSA	Financial Services Authority
FX	Foreign Exchange
GSTPA	Global Straight Through Processing Association
ICSD	International Central Securities Depository
ISIN	International Securities Identification Number
ISO	International Standards Organisation
ISP	Internet Service Provider
ISV	Independent Software Vendor
LAN	Local Area Network
LSE	London Stock Exchange
NBBO	National Best Bid and Offer
NMS	National Market System
NSCC	National Securities Clearing Corporation
NYSE	New York Stock Exchange
OPRA	Options Price Reporting Authority
OTC	Over The Counter
RSP	Retail Service Provider
SEC	Securities and Exchange Commission
SETS	Stock Exchange Electronic Trading System
SGML	Standard Generalized Markup Language
SIS	SegaInterSettle
SMS	Short Message Service
STP	Straight Through Processing
SWIFT	Society for Worldwide Interbank Financial Telecommunications
TOMS	Trade Order Management System
UTP	Unlisted Trading Privilege Plan
VPN	Virtual Private Network
VWABNBO	Volume Weighted Average Best Net Bid and Offer

VWAP	Volume Weighted Average Price
XML	eXtensible Markup Language

Appendix B

XML Protocols (Documents and Schemas)

This appendix contains a Purchase Order expressed as an XML document in Purchase order XML protocol (po.xml) and the XML schema defining Purchase Order XML protocol (po.xsd). This is an example reproduced from W3C (2001) , which we use to illustrate some of the germane features of XML Schemas.

Complex data types such as “USAddress” are easily defined in the Schema and can be incorporated in other complex data types such as “PurchaseOrderType”. Descendants can inherit the structure and definitions of their ancestors, much as they do in object-oriented programming.

The “PurchaseOrderType” contains an element of type “Items” (the list of items ordered) which is defined as containing between a preset minimum and maximum number of occurrences (defined in the Schema) of sub-elements of type “item” (the items ordered). Thus the schema constrains the validity of contents of a purchase order beyond the type of each element.

A “Stock-Keeping Unit (SKU)” data type is defined by the Schema and used as a “partNum” (part number) attribute of each “item”. The SKU is restricted to a special format of three numbers followed by “-“ and two letters and every item is required to be entered with an SKU which matches this pattern.

The XML parser of any system receiving a Purchase Order in this protocol will reject the Purchase Order if it violates the Schema in any way, for instance, if the PurchaseOrder contained larger numbers of “item” than the maximum allowed or if the SKU were missing or not of the required form. If the Schema is violated the Purchase Order XML document (the purchase order itself) is said to be “Not Valid”.

These and other features make XML schemas very flexible, extensible and powerful.

A Purchase Order Expressed in the Purchase Order XML Protocol (PO.XML)

```
<?xml version="1.0"?>
<purchaseOrder orderDate="1999-10-20">
    <shipTo country="US">
        <name>Alice Smith</name>
        <street>123 Maple Street</street>
        <city>Mill Valley</city>
        <state>CA</state>
        <zip>90952</zip>
    </shipTo>
    <billTo country="US">
        <name>Robert Smith</name>
        <street>8 Oak Avenue</street>
        <city>Old Town</city>
```

```
        <state>PA</state>
        <zip>95819</zip>
    </billTo>
    <comment>Hurry, my lawn is going wild!</comment>
    <items>
        <item partNum="872-AA">
            <productName>Lawnmower</productName>
            <quantity>1</quantity>
            <USPrice>148.95</USPrice>
            <comment>Confirm this is electric</comment>
        </item>
        <item partNum="926-AA">
            <productName>Baby Monitor</productName>
            <quantity>1</quantity>
            <USPrice>39.98</USPrice>
            <shipDate>1999-05-21</shipDate>
        </item>
    </items>
</purchaseOrder>
```

The purchase order consists of a main element, `purchaseOrder`, and the subelements `shipTo`, `billTo`, `comment`, and `items`. These subelements (except `comment`) in turn contain other subelements, and so on, until a subelement such as `USPrice` contains a number rather than any subelements. Elements that contain subelements or carry attributes are said to have complex types, whereas elements that contain numbers (and strings, and dates, etc.) but do not contain any subelements are said to have simple types. Some elements have attributes; attributes always have simple types.

The purchase order XML schema for the Purchase Order XML Protocol (po.xsd)

```
<xsd:schema xmlns:xsd="http://www.w3.org/2001/XMLSchema">

 <xsd:annotation>
  <xsd:documentation xml:lang="en">
   Purchase order schema for Example.com.
   Copyright 2000 Example.com. All rights reserved.
  </xsd:documentation>
 </xsd:annotation>

 <xsd:element name="purchaseOrder" type="PurchaseOrderType"/>

 <xsd:element name="comment" type="xsd:string"/>

 <xsd:complexType name="PurchaseOrderType">
  <xsd:sequence>
   <xsd:element name="shipTo" type="USAddress"/>
   <xsd:element name="billTo" type="USAddress"/>
   <xsd:element ref="comment" minOccurs="0"/>
   <xsd:element name="items"  type="Items"/>
  </xsd:sequence>
  <xsd:attribute name="orderDate" type="xsd:date"/>
 </xsd:complexType>

 <xsd:complexType name="USAddress">
  <xsd:sequence>
   <xsd:element name="name"   type="xsd:string"/>
   <xsd:element name="street" type="xsd:string"/>
```

```
    <xsd:element name="city"   type="xsd:string"/>
    <xsd:element name="state"  type="xsd:string"/>
    <xsd:element name="zip"    type="xsd:decimal"/>
   </xsd:sequence>
   <xsd:attribute name="country" type="xsd:NMTOKEN"
      fixed="US"/>
  </xsd:complexType>

  <xsd:complexType name="Items">
   <xsd:sequence>
    <xsd:element name="item" minOccurs="0" maxOccurs="unbounded">
     <xsd:complexType>
      <xsd:sequence>
       <xsd:element name="productName" type="xsd:string"/>
       <xsd:element name="quantity">
        <xsd:simpleType>
         <xsd:restriction base="xsd:positiveInteger">
          <xsd:maxExclusive value="100"/>
         </xsd:restriction>
        </xsd:simpleType>
       </xsd:element>
       <xsd:element name="USPrice"  type="xsd:decimal"/>
       <xsd:element ref="comment"   minOccurs="0"/>
       <xsd:element name="shipDate" type="xsd:date" minOccurs="0"/>
      </xsd:sequence>
      <xsd:attribute name="partNum" type="SKU" use="required"/>
     </xsd:complexType>
    </xsd:element>
   </xsd:sequence>
  </xsd:complexType>

  <!-- Stock Keeping Unit, a code for identifying products -->
  <xsd:simpleType name="SKU">
   <xsd:restriction base="xsd:string">
    <xsd:pattern value="\d{3}-[A-Z]{2}"/>
   </xsd:restriction>
  </xsd:simpleType>

 </xsd:schema>
```

The purchase order schema consists of a `schema` element, with the namespace set as an attribute, and a variety of subelements, most notably `element, complexType`, and `simpleType` which determine the appearance of elements and their content in instance documents.