10. Conclusions and Further Work

The main purpose for this thesis is an investigation into the systematic application of multirate DSP techniques to minimise the high computational bandwidth of classical AS in a TOB architecture. The rapid rise in CPU clock rates in recent years has meant that a software IFFT implementation of AS is perceived as rendering the custom hardware of a TOB obsolete. Simultaneously, interest in the TOB has declined because it is difficult to identify optimisations that do not compromise the generality of AS. However, it is argued that the chief functionality of AS is *the synthesis of musical notes* which have certain properties to which AS may be optimised, implying that TOB execution has high computational redundancy, thus leading naturally to the proposal of TOB optimisation by MAS. Unlike proposals for software IFFT synthesis, a noise model is excluded from MAS: justified on the basis that noise synthesis is not computationally intensive.

The idea of a time-invariant 'near-critical' sample rate is introduced for a time-varying oscillator x based upon *a priori* knowledge of its frequency trajectory during a finite lifecycle in note-based music synthesis. Exploitation of this property for sinusoidal oscillators in MAS creates two significant problems of (i) scheduling multiple sample rates and (ii) the interpolation of decimated sinusoids up to f_s to avoid quantisation noise. These are minimised by exploiting the *inherent hierarchy* of AS in that S sinusoids are summed into a final (multichannel) stream; the hierarchy is split and n_{fb} multirate filterbanks are introduced in the middle. A single filterbank interpolates many oscillators thus minimising interpolation overheads. Also, a small set of commensurate sample rates minimises scheduling overheads. The criteria for choosing a filterbank are (i) the 'goodness-of-fit' to expected oscillator distributions and (ii) how little the generality of AS is compromised. A binary-tree topology filterbank such as the QMF, when viewed as a *subband hierarchy*, is shown to provide a flexible and economic solution to both these desired aims.

Practical QMF implementation is shown to have two problems of (i) non-infinitesimal Δ_f and (ii) phase corruption through high-pass filtering. Economic FIR QMF designs have large Δ_f but are phase linear which is desirable for high functional transparency (to this end, efficient frequency and latency normalisation schemes are proposed). A suitable polyphase-allpass IIR QMF design has minimal Δ_f but is phase non-linear. High Δ_f is undesirable as it reduces the efficiency of a subband hierarchy because oscillators must be promoted into higher subbands when Δ_f 'deadbands' are removed by 'logical exclusion'. Using the properties of complex signal representation, a novel filterbank - the PEF - is introduced where Δ_f is 'physically excluded' and which is demonstrated in simulations to possess total phase transparency. As it is envisaged that QMF filterbanks will be implemented in software, the choice of FIR *versus* IIR can be left to the end-user, depending upon the required phase transparency. To its disadvantage, the PEF introduces significant extra complexity into the MAS algorithm.

A hypothetical scenario of the synthesis of a complete symphony orchestra provides two benchmarks - $E1$ and $E2$ - which quantify, respectively, how control bandwidth and computation (in multiplication frequency) are both consistently reduced across the range of filterbanks options. An important fact to emerge is that optimal QMF filterbank depth appears to be K=3, especially when FIR implementations are used. Supporting evidence is supplied when quantifying the performance of the proposed MASC in a real-time context parameterised by T_{max}. An upper bound of the 'speedup' offered by MAS is $E1=2^K$ for a stationary synthesis application mapped into the terminal subbands of the subband hierarchy. Filterbank computation, logical deadband exclusion and note non-stationarity reduce speedup, but a property of MAS is that computation is optimised to the note dynamics of the application in question and the generality of classical AS is preserved. Benchmarking is an area worthy of further investigation.

A survey of the two schools of sinusoidal oscillator design concludes that a phase-accumulator form is preferable to a recursive form for MASC implementation because of (i) the ease of dynamic frequency control and (ii) only one state variable is required in place of two for a 2^{nd} order recursive system. The commonly perceived bottleneck of sine transformation in a TOB via large LUTs can be solved using a pipeline form of linear-interpolated LUT where the interpolation is executed by a dedicated multiplier with an optimised wordlength. Another method is the CORDIC algorithm which can generate complex sinusoids in association with the PEF filterbank. Modifications are documented which result in the specification of an oscillator core based upon a CORDIC pipeline: a simulation shows how pipeline length and wordlength affect the output S/N. However, it is concluded that a pipelined linear-interpolated LUT provides the best solution for VLSI implementation.

The specification and design of the proposed VLSI MASC starts from a systems analysis of the processes in a hypothetical MAS synthesiser which form a synchronous logical pipeline. Data bandwidth proliferates in the MOB and is mapped into the MASC for computational acceleration; all other processes are mapped into software. The MASC is allied to a Shared Memory (SM) which is mapped via a bus interface into the Host CPU's address space to facilitate the required complexity of IPC. The Host has highest priority SM access and blocks the MASC, but as Host SM access is relatively infrequent compared to saturated MASC-SM traffic, MASC performance degradation is minimal. MASC functionality is encapsulated in the prototype Oscillator Descriptor which, in the context of an economical single data-bus MASC and coarse time-quantisation of PWL breakpoints, leads to the concept of concurrent MOB burst operation with interleaving of MASC-SM OD traffic. Dataflow in the MASC-SM system is thus organised on the principle of a *memory hierarchy* to optimise throughput.

Multirate scheduling and Host-MASC synchronisation are facilitated by a frame timing structure of length T_{frame} in which interleaved burst MOB operation is optimised for the lowest subbands at level K. Higher subband sample rates are generated by extending MOB bursts. Organising the round-robin schedule of OD's as a linked list facilitates (i) efficient ordering of OD's by subband and filterbank to enable, respectively, burst accumulation and mutually exclusive SM access to MOB-FBP buffers in the proposed synchronisation algorithm and (ii) efficient resource allocation by the Host when large blocks of OD's are manipulated. The logical process pipeline is driven by the FBP with which the MASC obeys a 'busy-waiting' protocol. Within frames, the Host and MASC operate with relative asynchronicity to enable the throughput of each to be maximised without the complexities of synchronous IPC in the context of a dynamically configured OD schedule. Finally, a dataflow diagram and the logic of the major functional units within the MASC are presented.

A fundamental fact to emerge from a parameterised analysis of MASC performance is that max(S) is a function of T_{max} and the MASC-SM databus bandwidth: to this extent the MASC is input / output bound and optimisation by MAS does not figure. The influence of MAS is to reduce the required internal MASC clock frequency from that required by a classical AS form of MASC by a factor of 2^{-K} which can be exploited by using cheaper, slower ASIC technologies: for K=3, internal MASC clock frequency is reduced eightfold. Assuming allocation to the deepest subbands at level K in a subband hierarchy, the same max(S) is achievable as in a classical AS form of MASC. However, a complexity increase arises in the MASC from supporting multirate operation by extended bursts due to the requirement for a variable length accumulation FIFO. Additionally, a software overhead is imposed by filterbank computation.

Future work can adopt a number of directions. Further refinements to the details of QMF subband hierarchy implementation are possible e.g. a deeper analysis of optimal K in the context of comprehensive benchmarking. Also, the development of a MASC behavioural simulation and prototype will permit a more detailed analysis of actual real-time performance than that anticipated at the level of a functional specification. Another area which is unexplored is the use of MAS as a software synthesis technique. However, a final comment upon this thesis is that the proposed MAS algorithm and consequent MASC architecture are, in the opinion of the author, about as far as TOB-based AS implementations can be realistically optimised: a significant improvement over the classical TOB is feasible. So this work may be comprehended as both (i) a necessary summary to a line of research that has fallen out of favour before achieving its natural level of maturity and (ii) a potential starting point for a revival of interest in TOB methods.