Very important: ensure that you provide answers for every box on the reviewer's form. A copy of the standard responsive mode form is at the end of this document.

Do ensure that you copy most of the information over to the associated EPSRC form. Don't worry about duplicated information.

Research Proposal

Communication Centric Microelectronic Design

Simon Moore and Robert Mullins
Computer Laboratory
University of Cambridge
16th May 2005

Make it clear why your team is highly capable of undertaking the proposed research

Part 1: Previous Research and Track Record

1 Microelectronic Design at University of Cambridge

The Computer Laboratory at the University of Cambridge continues to achieve top ranking in research assessment exercises. Since starting as a lecturer in 1998, Dr Simon Moore has been building a group undertaking research in microelectronic design with particular emphasis on self-timed circuits, embedded processor architecture and secure hardware. Three of his PhD student graduated last year and he currently supervises nine PhD students and has three RAs. The group not only advances the theory of the subject, but continues to build prototype chips which provides a healthy dose of reality. Whilst we find chip design to be a very important component, it requires a great deal of support and corporate memory. The group, at 13 people, has reached critical mass. Dr Robert Mullins (coinvestigator to be appointed as a senior RA on this project) has been pivotal in creating design flows which fit our research needs as well as playing a key rôle in architectural- and circuit-level design.

2 Principle Investigator

Dr Simon Moore (Senior Lecturer) leads the microelectronic design group. He has an MEng from University of York (first class) and a PhD from University of Cambridge on multithreaded processor design. His thesis work was published as a book by Kluwer [12]. He spent a summer working at DEC Western Research Lab. in Palo Alto. After his PhD he worked as an RA in Cambridge advancing self-timed circuit techniques (funded by an EPSRC grant GR/J62708, 1994–1997 and AT&T in 1998).

In 1998 he was appointed as a lecturer in Cambridge and begun growing a group looking at self-timed circuits and embedded processor architecture. This started (1998-2001) with a project to advance and evaluate self-timed circuits for embedded processors funded by EPSRC (grant GR/L86326/01) in collaboration with Cambridge Consultants and AT&T Research. This lead onto the fusion of clocked and self-timed techniques as part of the GALS (globally asynchronous, locally synchronous) project (EPSRC grant GR/R52299/01, 2001-2004) in collaboration with University of Manchester and Marconi. These techniques were also advanced as part of a collaborative project with MIT looking at networks on chip for tiled processor architectures (CMI, Jan 2003 to Dec 2005). Work on building robust asynchronous circuits for smart cards has also provided a great deal of insight into novel circuit techniques, initially with EU funding on the G3Card project (IST-1999-13515, 2000-2003) and later through a collaboration and funding with Epson and EPSRC (GR/S42897/01, 2003-2006). Projects on wearable embedded electronics (Microsoft, 2004-2007) and coding schemes for multiwavelength optical networks (Intel, 2004-2007) complete the current portfolio of work. He is currently a Visiting Scholar at Intel.

Dr Moore has over 45 publications and has won a pest paper award at Async'2002. He has given invited talks at eSmart (Cannes, 2001), IEEE International Solid States Circuits Conference (San Francisco, 2002) and the Royal Academy of Engineering (London, 2003). He presented a tutorial on "Breaking the synchronous barrier for systems-on-chip" at DATE (Paris, 2004). He holds two patents (now both covered in multiple countries): one on a microprocessor resistant to power analysis, and the other on a sensitive capacitance sensor for biosensor applications.

3 Coinvestigator

We plan to employ Dr Robert Mullins as a Senior RA on this project. Robert obtained BEng, MSc and PhD in parallel asynchronous processors from Edinburgh University before starting as an RA in Cambridge. He has played a key rôle in designing three test chips in Cambridge, the most recent of which was returned in early March 2005 code named Lochside (Figure 1). The Lochside test chip includes a high performance low latency network on chip which was designed in collaboration with MIT (funded by CMI) and University of Manchester (funded by EPSRC grant GR/R52299/01).

The architecture was presented at the prestigious ISCA conference last year [13] and we anticipate several publications based on measured results from the test chip. The network-on-chip architecture we developed was tailored for a tiled processor architecture (SCALE) being developed by our collaborators at MIT. However, it became apparent that this network-on-chip has far wider application. In particular, as we move into the nano CMOS era, wires dominate performance in terms of delay and power consumed. Consequently we believe that microelectronic design must become far more communication centric, with a range of networks-on-chip providing a critical component.

Unnecessary random picture! You could fill up this area with references to your related work.

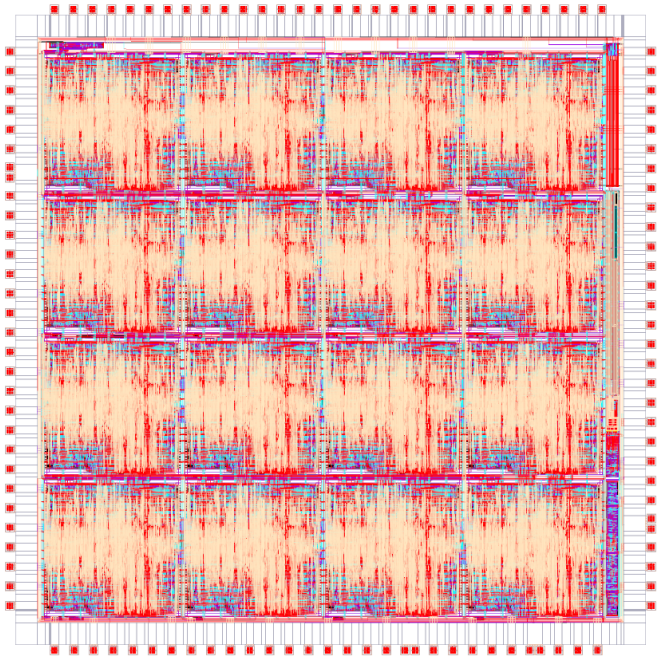

Figure 1: Lochside test chip which includes a network-on-chip

Grab the reviewer early. State the research project in the first paragraph.

Don't cram too many words onto a page (e.g. font that is too small or lines that are too long) since it will make the proposal hard to read which may annoy the reviewer.

Part 2: Proposal

4 Background

We wish to undertake research into communication centric microelectronic design methods which are in contrast to today's computation centric (or *gate-level*) techniques. We believe that this research is timely since electronics is at the cusp of change. For the last 60 years, digital gates have been costly to produce and have limited performance. We are now entering an era where the wires, which were once almost free, becoming the cost and performance limiter. This trend is well documented in the International Technology Roadmap for Semiconductors (ITRS) roadmap which clearly identifies the step change needed in circuits and associated design techniques. They also identify spiralling design complexity, escalating power densities and associated thermal problems.

Then make it clear why the research is required.

4.1 Challenges

The research proposed here is driven by five observations:

Wire delay — Wires, rather than gates, are increasingly the speed limiter in IC designs targeting deep submicron CMOS. Over the years, chip area has remained almost constant for many devices (microprocessors, FPGAs, etc.), and yet wire delay per unit length continues to increase. Communication across chip requires repeater buffers at the very least. Replacing buffers with flip-flops allows the wires to be pipelined, thereby improving utilisation. This observation leads one to think about other architectural possibilities.

Clocking — Clocks are typically used to sequence operations on chips. Global clock distribution with low skew to all nodes in the circuit is untenable in the future. Our research into distributed clock generation and distribution [3] indicates that time can be uniformly distributed across future deep submicron chips. None-the-less, even with advanced clocking techniques it is still important that long distance data traversal needs to take into account the limitations of the clock distribution technique used.

Signal integrity — As wires have become narrower, they have also become taller in order to reduce resistance per unit length. Unfortunately tall wires placed close together couple rather well which at best introduces data dependent timing characteristics and at worst corrupts the data. Consequently, signal integrity is an increasing issue for longer wires. Therefore, there is increased engineering effort in the layout and analysis of long wires.

Cross sectional bandwidth — Rent's rule [15] indicates that the cross sectional bandwidth of a design inexorably increases which has resulted in an ever increasing demand for layers of metallisation to provide sufficient static routing capacity. This in turn has an impact on mask costs for chip manufacture which are currently growing exponentially (currently around $1m for a mask set for a 90nm process).

Design for reuse — Ever increasing design complexity is being tackled using the well trodden path of abstractions or *modularisation*. Well defined modules can then be designed for reuse, often referred to as *intellectual property blocks (IP blocks)* in the EDA business. Clean interfaces are required for quick and reliable integration. Commercial examples of communication interfaces include ARM Ltd's AMBA bus for ASICs and Altera's Avalon bus for FPGAs.

4.2 The communication centric design approach

Explain the research approach to be taken.

From the five challenges it is apparent that movement of data across chip is an increasingly important issue. **Based on our research [13, 11][1] and that of others [1, 4], we believe that the design of effective on-chip networks will be key to achieving both power and performance goals and**

[1]Funded through two projects, one by EPSRC (GR/R52299/01) and the other by the Cambridge MIT Institute (CMI)

alleviating the problems of integration and verification in future ASICs and FPGAs. Such an interconnect can be designed for reuse. Top level wiring can be well designed to ensure high performance and good signal integrity. Routing data through switches allows wire reuse. Wiring between switches can be routed with geometric regularity which simplifies verification of signal integrity. But what architectural form should the networks take for different applications and how can designs be technology mapped onto FPGAs and ASICs?

4.3 Design metaphors

Viewing a switched interconnect as a bus replacement is just one design metaphor. One could also provide an abstraction which looks like independent FIFOs or communication channels. This abstraction maps well to constructs appearing in hardware description languages (HDLs) like SystemC and SystemVerilog. For existing design systems, generators like Altera's SOPC builder could be used to instantiate FIFOs, communication channels and bus equivalents that would map directly onto the switched interconnect.

Dynamic scheduling of the switches may well be appropriate for many applications, but others may require hard real-time requirements to be met. One option is to statically schedule the switches, thereby time division multiplexing the wires.

The latency of chip wide communications will require architectural changes, some of which can be addressed through high-level language support which enable latency tolerant communication to be described. Once a design is represented in this abstract form, constraint-driven communication synthesis can then be performed [14].

4.4 Technology mapping

Higher level HDLs and reuse of abstract IP blocks allows a design team to construct complex systems more rapidly. This design approach typically allows designs to be targeted at a range of implementation technologies, from soft programmed devices like FPGAs, to gate array type structures (e.g. Altera's *Hardcopy* devices) and hard wired ASICs. A switched interconnect could present the same interface for all levels of technology mapping, though the implementation may be quite different:

FPGA — On an FPGA the interconnect will augment some of the longer statically configured wires. Some full custom components may be required, for example custom drivers to allow high speed time division multiplexing on wires. Such multiplexing could then run at least $10\times$ faster than the surrounding programmable FPGA fabric. For the rest of the architecture it is currently unclear which other components (e.g. a cross bar, arbiters, etc.) are best implemented in programmable logic and which will benefit from full custom implementation. It may also be beneficial to co-locate full custom components with specific reconfigurable components. So, for example, one or more embedded memories might be co-located with the NOC core, thereby allowing the amount of buffering to be configurable, or to provide routing tables for more complex architectures.

Hardcopy — Some critical components of the switched interconnect could be hard coded into modern gate array technologies (e.g. Altera's Hardcopy devices).

ASIC — The switched interconnect could be laid down as a full-custom or semi-custom generated structure as part of a block based design flow.

4.5 Related work

Once you've explained your research idea, put it in context (i.e., avoid putting too much background too early which can detract from your research idea).

Commercial bus systems are already evolving into networks-on-chip in order to address today's implementation technology demands. For example ARM's multilayer AHB, IBM's CoreConnect, ST Microelectronics' STBus, Hitachi's SuperHyway, Sonics' SMART interconnect, siliconbackplane and Manchester's CHAIN interconnect, etc. These all replace the bus with some form of cross bar with central arbitration, neither of which scales well to a large number of ports.

Research into networks-on-chip has produced several theoretical architectures, but few have been implemented. Notable exceptions are the static and dynamic networks of MIT's RAW tiled processor architecture [5, 6], Philips' SOC interconnect [2] and of course our own work [13].

4.6 Originality

Make it clear why your proposal contains an original research idea. Don't leave it up to the reviewer to think/guess!

Our background work into networks-on-chip has already made advances at an architectural level using speculative precomputation of control [13]. This technique allows the control to be moved away from the critical path, thereby reducing latency. Where as existing designs take three or more cycles to traverse each router node, our current design takes just one cycle without impacting the cycle time. With the architectural and VLSI design expertise we already have we are in an excellent position to push the architecture further, looking at a wider range of applications and technology mappings. Where others have undertaken architectural level research into network-on-chip structures, we are in the rare position to have fabricated a design (chips returned and tested in March 2005 with a paper submitted to ICCD containing the results). This gives us significant insight into the implementation complexities.

We also have experience in designing complex systems. For example, an earlier test chip of ours (Springbank [7]) contained five processors using a range of synchronous and asynchronous design styles, a globally asynchronous locally synchronous communications network [9, 8, 10], an I/O block and 2048-bit Montgomery multiplier to accelerate RSA public-key cryptography.

5 Programme and Methodology

5.1 Aims

Clearly state the aims of the research which draws together the material discussed in the Background section.

We would like to undertake a 3 year research project into *communication centric microelectronic design* in the Computer Laboratory at the University of Cambridge in collaboration with Altera (UK). This work will move us away from the traditional computation or gate level design style toward a communication centric design style. Network-on-chip structures will provide the core mechanism for across chip communication in this new design style. The network could be quite simple for more local or statically scheduled communication, through to more elaborate dynamically scheduled communication maximising chip wide bandwidth. We envisage that the network will not only provide a communication infrastructure, but will also play a key rôle in scheduling resources and power management.

5.2 Objectives

Clearly state the objectives.

To model, design and fabricate network-on-chip architectures suitable for embedded systems design. Full custom networks for both ASICs and FPGAs will be addressed since both of these target technologies are of commercial interest and yet provide different design constraints.

This work can be broken down into the following more detailed objectives:

1. Compile a set of benchmark applications and technology parameters with assistance from our industrial and academic partners.
2. Create a highly parameterised model of a network-on-chip to allow rapid architectural exploration.
3. Perform the technology mapping from the architectural model to transistor level implementations for both ASICs and FPGAs.

5.3 Programme of Research

How will the research be undertaken and explain who will do what.

The programme of work will be undertaken by two PhD students and a Senior RA. One PhD student will specialise on the ASIC target for the network-on-chip, and the other PhD student will specialise on the FPGA target. A Senior RA (coinvestigator) and PI will provide the background experience for all aspects of the project. In particular, we already have one architectural model of a network on chip which was presented at the prestigious ISCA conference [13], and we have completed a test chip with implements an improved variant of this architecture (chip returned in late March 2005 with results send to the ICCD conference). This work introduced a novel architecture which reduced communication latency using speculative precomputation

control. Whilst we found our initial architectural modelling to be hugely valuable, addressing the implementation nuances made a significant impact on the detailed architectural decisions.

From this background and the objectives stated earlier, the following programme of work follows:

Benchmarks — Identify application areas and how they can be mapped to a network-on-chip style system with assistance from our industrial partner, Altera. We will also build on our experience of implementing cryptographic coprocessors, handling streamed media, etc.

Architectural Modelling — Replace our existing architectural model of a network-on-chip to produce a more generic and parameterisable model to allow greater exploration of the architectural possibilities.

ASIC target — Implement key components of the network-on-chip architecture. This data will then allow refinement of the architectural model and help us better understand how efficiently applications can be mapped to this technology.

FPGA target — Look at the details of a network-on-chip comprising reconfigurable and full custom components for future generation of FPGA. This work will leverage results from the ASIC implementation but will also need to address the coupling to existing FPGA components (embedded memories, look-up tables (LUTs) and programmable wiring). We will also explore how the presence of such a network can be exploited when mapping applications to this new FPGA architecture.

5.4 Work Packages and Personnel

Details of deliverables and how the work will be allocated.

It is envisaged that three people will work full time on this project for three years. Dr Robert Mullins will be appointed at a Senior RA level. He will provide the central thrust and focus for the project in conjunction with the Principle Investigator. Two project based PhD students will work on the architecture and implementation for ASIC and FPGA version of networks-on-chip. They will share the common goal of establishing a benchmark suit and architectural simulation base. This work is broken down into the following work packages:

WP1 Literature survey and first year reports from PhD students.

WP2 Produce a benchmark suite (all collaborate on this).

WP3 Highly parameterisable network-on-chip simulation model (lead by SRA).

WP4 Research FPGA network-on-chip architectures and technology mappings (one PhD student with architectural input from SRA).

WP5 Research ASIC network-on-chip architectures including wire optimisation, gate-level design and some full custom components. Publish papers based on simulation results (one PhD student with architectural and VLSI design input from SRA)..

WP6 Based on winning ASIC design simulations, take the winning designs through to layout and ship test chip. Publish papers based on simulation results (all collaborate on this).

WP7 Evaluation report for the test chip. Publish papers based on implementation results. All collaborate on this.

WP8 Produce two PhD thesies (PhD students).

In addition to the work packages, academic papers will be produced during the project as results become available. Quarterly meetings will be held with our UK industrial partner and annual meetings with our MIT partner with email and phone conferences in between.

5.5 Management of Risk

The reviewer has to comment on the level of risk, so make it clear that it is a bit risky (i.e., interesting) but it is not dangerous (i.e., the work will never be completed).

The principle risk is managing the complexity of the project. Benchmarks are required to give us real traffic flows rather than the current simple statistical models we use at present. The architectural modelling needs to be firmly grounded in what is implementable. But in all of these areas, the investigators have a strong track record of meeting similar challenges on other projects including the Lochside test chip already mentioned, and the Springbank test chip [11] which included five processors each designed using different clocked or asynchronous techniques.

Research into networks for future FPGAs is a departure from our current experience, but our experience of ASIC design flows and our collaboration with Altera should provide ample expertise.

5.6 Collaboration

The reviewer must comment on collaboration, so make it clear what collaboration is likely to happen, preferably with letters of support.

Industrial: The principle industrial collaborator on this project is Altera (UK). They have worked with us over the years on a number of undergraduate teaching projects, and have fully funded one PhD student (Dr Frankau) in the past (Dr Frankau's thesis was recently nominated for the BCS distinguished dissertation award). **Please find the letter of support attached.**

Academic: Out network-on-chip research has been in collaboration with Prof Krste Asanovic at MIT funded by CMI. This project is due to end in December 2005 and CMI currently has no funding available to extend the project (CMI's DTI funding is due to finished in 2006). We very much hope and plan to continue collaboration with Prof Asanovic, but joint funding is proving to be problematic. Consequently this research proposal stands in its own right and Prof Asanovic is bidding for funding in the USA to undertake closely related research. **Please find the letter of support from Prof Asanovic attached.**

We have an on going collaboration with Prof Furber's group at the University of Manchester. They are funded from an EPSRC portfolio grant so cannot be part of this application, though we have received joint funding from EPSRC and the EU in the past. For example, the Lochside test chip returned in March 2005 included design from Manchester.

5.7 Dissemination and Exploitation

Do more than say "we'll publish a few conference and journal papers".

We have a track record in publishing architectural level papers in international conferences like ISCA and ICCD, and more circuit oriented papers at Async. We have presented tutorials and invited talks at DATE and ISSCC, and believe that these conferences will be good targets for implementation level results.

Early on in the project we will need to produce a benchmark suit which we will make publicly accessible to allow others to use the benchmarks and also to encourage open scrutiny of the work. This, and other aspects of the research, will benefit from close links with our industrial and academic partners.

We have hosted a number of UK Async Forums and attend a wide range of workshops around Europe. We ran a four day residential school on *Timing for Deep Submicron Chips* in January 2005 which was very successful, and hope to offer similar events in the future. Building up infrastructure and a common vision for research in microelectronic design in the UK is a passion of ours, and we will continue to fully contribute to events and activities in this area.

5.8 Justification of Resources

Justify resources requested. In my experience, all resources that have a modicum of justification are awarded.

Should we be doing more to bring in funds for general Lab equipment infrastructure? Shopping list from sys-admin's?

We request resources for three full time personnel:

1. One PhD student to address the network-on-chip architecture for ASICs
2. One PhD student to address the network-on-chip architecture for FPGAs
3. One Senior RA (Dr Robert Mullins) to act as principle architect for the project. Dr Mullins is rapidly building an international reputation in networks-on-chip, and is in an excellent position to fulfil this rôle (see Part 1). We request that sufficient funds are granted to allow one discretionary point increase since recruitment and retention of staff is difficult given the buoyant local technical economy.

We also request enough resources to fabricate one test chip on a 0.13μm process through Europractise, together with the PCs and ECAD tools to support this endeavour. Travel resources are requested to attend conferences and visit our collaborative partners. 10% of a computer officer is requested to help maintain the ECAD tool installation (we currently have 20GB of installed tools).

References

[1] W. J. Dally and B. Towles. Route Packets, Not Wires: On-Chip Interconnection Networks. In *Proceedings of the 38th Design Automation Conference (DAC)*, June 2001.

[2] J. Dielissen, A. Rădulescu, K. Goossens, and E. Rijpkema. Concepts and implementation of the Philips network-on-chip. In *IP-Based SOC Design*, Nov. 2003.

[3] S. Fairbanks and S. Moore. Self-timed circuitry for global clocking. In *Eleventh International Symposium on Advanced Research in Asynchronous Circuits and Systems*, New York, 2002.

[4] A. Jantsch and H. Tenhunen, editors. *Networks on Chip*. Kluwer Academic Publishers, 2003.

[5] M. B. Taylor *et al.* The Raw Microprocessor: A Computational Fabric for Software Circuits and General Purpose Programs. *IEEE Micro*, March/April 2002.

[6] M. B. Taylor *et al.* Evaluation of the raw microprocessor: An exposed-wire-delay architecture for ilp and streams. In *Proceedings of the 31st Annual International Symposium on Computer Architecture (ISCA)*, 2004.

[7] S. Moore, R. Mullins, and G. Taylor. Designing one-of-four encoded datapaths. In *12th UK Async. Forum*, June 2002.

[8] S. Moore, R. Mullins, and G. Taylor. The Springbank test chip. In *12th UK Async. Forum*, June 2002.

[9] S. Moore, R. Mullins, and G. Taylor. Simple switched GALS interconnect. In *Third ACiD-WG Workshop (5th Framework Programme), FORTH, Crete*, Jan. 2003.

[10] S. Moore, G. Taylor, R. Mullins, and P. Robinson. Channel communication between independent clock domains. In *Fifth ACiD-WG Workshop, Nauchatel*, 2001.

[11] S. Moore, G. Taylor, R. Mullins, and P. Robinson. Point-to-point GALS interconnect. In *Eighth International Symposium on Advanced Research in Asynchronous Circuits and Systems*, Manchester, England, 2002.

[12] S. W. Moore. *Multithreaded Processor Design*. Kluwer Academic Press, ISBN 0-7923-9718-5, Apr. 1996.

[13] R. D. Mullins, A. F. West, and S. W. Moore. Low-Latency Virtual-Channel Routers for On-Chip Networks. In *Proceedings of the 31st Annual International Symposium on Computer Architecture (ISCA)*, 2004.

[14] A. Pinto, L. Carloni, and A. Sangiovanni-Vincentelli. Constraint-driven communication synthesis. In *Design Automation Conference (DAC)*, 2002.

[15] A. Tetelbaum. Generalizations of rent's rule. In *27th Southeastern Symposium on System Theory*, 1995.

Overview of the Programme of Research

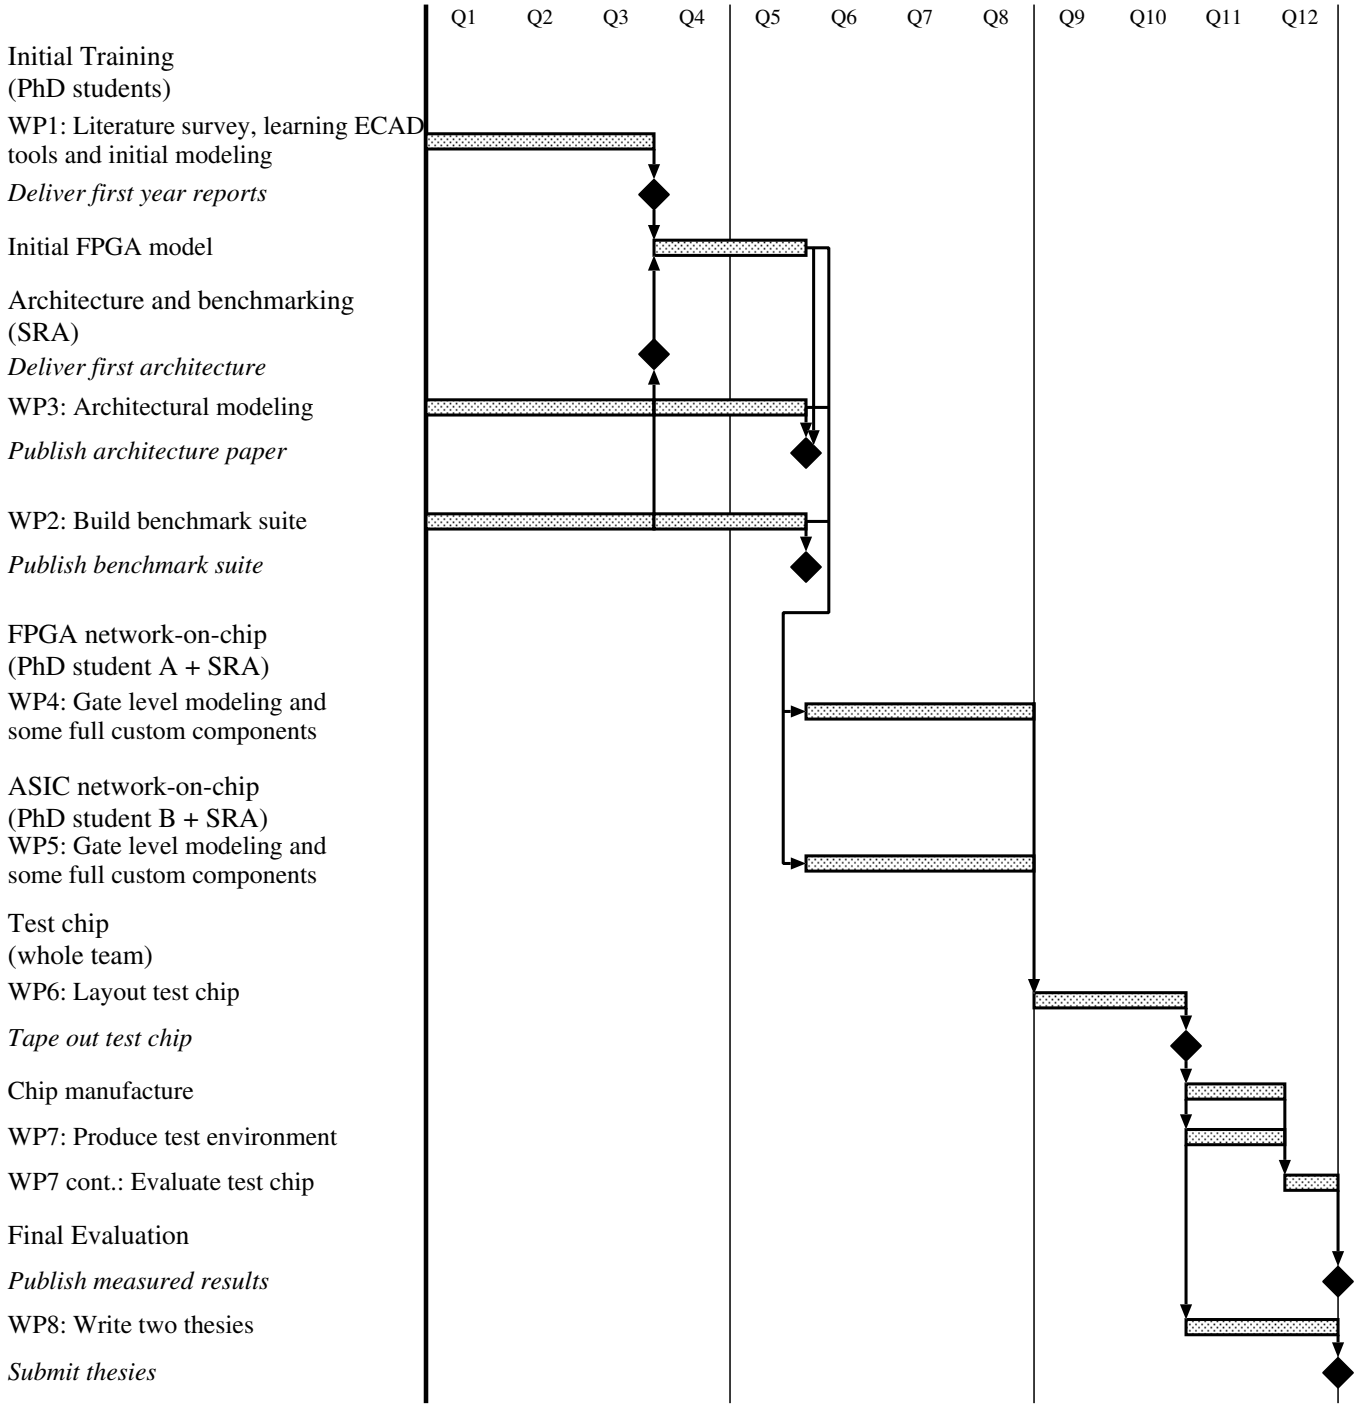

Colaborative Exchange

Meetings with Altera and MIT will take place at least every quarter.

This is the sort of form that reviewers (typically academics like yourself) have to complete. Make sure that it is easy for the reviewer to fill in all of the boxes, particularly for the sections over the page like "dissemination" and "resource management" which are scored and are, therefore, subject to simple analysis (i.e., how many outstanding/major boxes are ticked).

STANDARD RESEARCH PROPOSAL

REFEREE'S ASSESSMENT

1. Proposal Details	Proposal Reference No(s)

Important: If you feel you may have a **conflict of interest**, or if you have a **low level of confidence** in your ability to provide an assessment of this proposal**, please advise EPSRC before proceeding**.
In the case of a strongly multi- or interdisciplinary proposal, please comment only on the elements of the proposal within your area of expertise. Please say what this is in Section 4.

Title: Investigator
Organisation

In case of any query relating to this assessment contact and return your assessment of this proposal to them by:

Referee reference:

2. Your Assessment Please provide a narrative assessment addressing the following criteria and place a cross in the appropriate boxes.

Research Quality Comment on the quality of the proposed research.	
Methodology Comment on the feasibility of the proposed methodology.	
Adventure in Research Comment on the research risk and originality in this proposal and its balance with the potential to make major advances in knowledge. Where the proposal is more incremental in nature, comment on the degree to which this is essential to the advancement of knowledge in the field.	
The overall **research quality** of this proposal is: **Unsatisfactory** ☐ **Adequate** ☐ **Good** ☐ **Outstanding** ☐	
My confidence level in assessing this is: **Low** ☐ **Medium** ☐ **High** ☐	
Research impact Comment on the extent to which the results/new knowledge likely to accrue from this research will impact upon The academic community The industrial/non-academic sector The overall likely **impact** of this work will be: **Low** ☐ **Limited** ☐ **Significant** ☐ **Major** ☐	
My confidence level in assessing this is: **Low** ☐ **Medium** ☐ **High** ☐	
Academic Collaboration Where collaboration is proposed with another department or Institution, comment on the likely value associated with the partnership.	
Non-Academic Collaboration Where industrial or non academic collaboration is proposed, comment on the strength of the partnership and the degree to which appropriate contributions to the research project are to be made.	

Dissemination Comment on the degree to which the arrangements for dissemination are likely to be effective in informing the academic and, where appropriate, the non-academic community.	

The overall **collaboration and dissemination** arrangements are: **Unsatisfactory** ☐ **Adequate** ☐ **Good** ☐ **Outstanding** ☐	

Resources and Management ***Planning and Project Management*** Have appropriate project management plans been specified, including realistic timescales and milestones?	

Resources Requested Do the research objectives justify the requested resources?	

The overall **management arrangements** of this project in order to fulfil its objectives appear: **Unsatisfactory** ☐ **Adequate** ☐ **Good** ☐ **Outstanding** ☐	

My confidence level in assessing this is: **Low** ☐ **Medium** ☐ **High** ☐	

3. Your Conclusions Please use the space below to make any comments you may wish, relating to this proposal, not made elsewhere.

Overall the quality of this proposal as written is: **Unsatisfactory** ☐ **Adequate** ☐ **Good** ☐ **Outstanding** ☐	I believe this project: **Should not proceed as proposed** ☐ **Could proceed as proposed** ☐ **Should proceed** ☐

4. Your Area of Expertise Please indicate what your area of expertise is that is relevant to your assessment of the proposal.