Ref. No.-ITS/027/07 Date/Time of Receipt-   (For Official Use)

Application for
the Innovation and Technology Fund

Innovation and Technology Support Programme

Normal Tranche (Tier 3)

A Compiler for High Performance Computing on Array Technologies

Submitted by
The Chinese University of Hong Kong

Section A – Basic Information

(When an application is approved, the information in this section will be published at the ITF website (denoted with #) and/or carried forward as pre-filled data in the related progress/final reports.)

1. Project Title#

1.1. English (not more than 20 words):

A Compiler for High Performance Computing on Array Technologies

1.2. Chinese (not more than 30 words):

陣列技術高性能計算編譯器

2. Abstract#

(Please provide a brief summary of the objectives, R&D methodology and the impacts and benefits of the project.)

2.1. English (not more than 200 words):

The new multi-core, reconfigurable computing and graphics processing unit technologies, referred to as "array technologies" (AT), offer improved price and performance over conventional high performance computers (HPC). They can often sustain higher numbers of floating point operations per second as they are more specialized than traditional processors, and are widely regarded as an important direction for HPCs. ATs are particularly suitable for industries in Hong Kong and Mainland China as air conditioning and machine room space constraints are of utmost importance.

The main barrier for the adoption of array technologies is that even for experts, they are difficult to program. Since decades of compiler research has not resulted in a general solution, we limit our study to an important class of problems that include Monte Carlo Simulation (MCS) and optimization, and propose to develop a self-checking, self-tuning, source-to-source compiler for the utilization of ATs. The resulting tool can be used in applications including derivative pricing, risk management, portfolio optimization, time series prediction and enterprise resource management, all of which are industries of great importance to Hong Kong. The solution will lower the barrier for entry both in terms of price-performance and user-friendliness and enable the widespread use of array technologies.

2.2. Chinese (not more than 300 words):

新一代的“陣列技術(AT)”，包括多核、可重構計算、以及圖形處理器(GPU)幾種技術，大幅提高了傳統高性能計算機(HPC)的性能價格比。這些技術能支持更高速的浮點運算，是公認的高性能計算的重要發展方向。陣列技術尤其適合香港和中國大陸的工商業界：陣列技術能大幅降低計算系統的降温耗電和所佔空間。

推廣陣列技術的主要障碍是在其編程困難。科研界經過數十年的努力仍然未能找出一種可以普遍通用的辦法使编程變得容易。在本研究計劃，我們建議針對香港業界的需要，發展一套專門爲 Monte Carlo 模擬及優化計算的陣列技術解决方案，其核心是一套能進行自我檢驗、自我調校、源到源的編譯器。Monte Carlo 模擬及
優化計算廣泛應用於香港工商業界，包括衍生工具定價、風險管理、投資組合優化、時序預測、企業資源管理等。

我們的解决方案將在提高性能價格比和易於使用兩方面，讓香港工商業界廣泛使用新一代的陣列技術。

3. Objectives

(Please state in clear and specific terms the technological challenges and what the proposal aims to achieve.)

1. High level development. Compiler researchers have unsuccessfully tried to develop automatic, parallelizing compilers for many decades. We believe that the general problem will not be solved in the near future. However, with array technologies becoming pervasive, some simple programming method is required. By restricting our applications to narrow domain, we aim to develop a workable parallelizing compiler for the different array technologies studied.
2. Efficiency. It is difficult to generate code for modern architectures, particularly those that vary greatly in nature such as for different array technologies. Nevertheless, we aim to produce highly efficient code in our given domains.
3. Accuracy. Array technologies can perform single precision computation much faster than double precision. Furthermore, mathematics libraries that can tolerate small inaccuracies run much faster than strict ones. Our compiler should be able to generate self-checking code that operates in single precision but gives a warning if the results are likely to be inaccurate.
4. Portability. We aim to develop the first compiler that takes a standard C program as input and can generate efficient and accurate implementations on multi-core, GPU and FPGA platforms, making it possible for the first time to write portable code for all of these technologies. Moreover, it is possible to write a program and then execute on the different platforms to determine the most suitable one for a given application.
5. Market share. Our tool should be able to portably and efficiently compile standard C programs in our domain to the three different platforms (reconfigurable computing, GPU and multi-core). This will allow our users to easily deploy array technologies in production computing environments. We aim to develop the market for this type of domain-specific parallelizing compiler for array technologies and capture a large market share.

4. R&D Methodology

(Please provide a brief description of the technology to be developed and/or the innovative use of existing technologies. Details should be provided in Section C.)

The five objectives will be respectively addressed as follows:
1. High level development. We will specialize our compiler to applications in optimization and Monte Carlo Simulation (MCS) which are used in financial derivative pricing, risk management, time series prediction, enterprise resource management and traffic simulation.
2. Efficiency. We will use automatic tuning techniques in which our compiler first generates different implementations of a program with various parameterized optimization techniques. The programs are then executed with different parameter settings, and the parameters are automatically adjusted in order to minimize execution time.
3. Accuracy. We will develop techniques that perform computations in single precision and check the results in double precision. If the results are not consistent, a warning will be reported and the entire application run in double precision to ensure that it degrades gracefully.
4. Portability. Our compiler will have different back-end code generators which can target the three different types of array technology studied.
5. Market share. We believe that as multi-core technologies become pervasive, our approach will be technically superior to more general ones. By applying the ideas above, together with marketing help from Cluster Technology, we believe that we can develop and be a leader in this market.

5. Deliverables#

(Please set out the deliverables, itemized as appropriate, together with their *detailed technical specifications*, of the project.)

5.1. English

1. Efficient implementations of MCS and optimization routines on the platforms studied. Case studies for a number of applications (MCS pricing of a financial derivative, grid search, simulated annealing and genetic algorithm) will be made.
2. Techniques for generating self-tuning code that can adjust optimization parameters and deduce the best settings for a given hardware platform.
3. Techniques to generate self-checking routines to verify the accuracy of single precision MCS runs using a small number of double precision runs.

4. A source-to-source compiler that can take a C language program in our specified domains (MCS and optimization) and produce efficient and accurate implementations on the three array technologies studied.
5. A website and papers that serve as a tutorial for the compiler as well as describe our compiler and results of applying the infrastructure developed to solve practical problems in derivative pricing.

5.2. Chinese

1. Monte Carlo 模擬及優化計算在陣列技術平台上的有效實施，包括衍生工具定價,網格搜索，退火模擬及基因算法幾種案例研究.
2. 發展自我調校技術，能夠自動產生軟件，以針對硬件配置選擇最佳參數值.
3. 發展自我檢驗技術，能夠自動產生軟件，以小量雙精度計算核實單精度 Monte Carlo 模擬的準確度.
4. 發展一個源到源的編譯器，可以在多核、可重構計算、以及圖形處理器等三種陣列技術平台上有效地實施 Monte Carlo 模擬及優化計算的 C 编碼。
5. 以論文及網頁形式介紹上述編譯器及其在衍生工具定價的實際應用結果。

6. Dissemination Plan

(Please provide a brief description of the activities that would be undertaken to promote the project deliverables to the target users. Details should be provided in Section C.)

The main avenues for dissemination are international conferences, via a project website and user forum. Documentation in both Chinese and English will be provided. We will also take advantage of Clustertech's existing client base and directly approach suitable organizations for related project and consulting work.

Section B – Implementation Schedule

(When an application is approved, the information in this section will be published at the ITF website (denoted with #) and/or carried forward as pre-filled data in the related progress/final reports.)

1. Overall Schedule#

(When planning the project commencement date, please consider that it may take of a few months to approve an application.)

Commencement date (dd/mm/yyyy) : 1/10/2007

Completion date (dd/mm/yyyy) : 30/9/2008

Project duration (month) : 12

2. Project Milestones#

(Please set out individual milestones to be achieved at different stages of implementation. Detailed account should be provided Section C.)

Period		Milestones
From (dd/mm/yyyy)	To (dd/mm/yyyy)	
1/10/2007	15/11/2007	Hardware evaluation: Evaluate different hardware platforms and software development kits. Purchase hardware. Produce project website with initial project information which will be updated throughout the project.
1/11/2007	30/11/2007	Libraries and Benchmarks: Development of libraries for uniform and Gaussian random number generators, optimization routines, primitive operations etc for multi-core, FPGA and GPU technologies. Development of benchmark applications.
1/12/2007	30/4/2008	Compiler: Development of compiler to translate MCS and optimization programs to multi-core, FPGA and GPU implementations.
1/5/2008	30/9/2008	Compiler: Development of features for self-tuning code generation and self-checking of result accuracy.
1/7/2008	30/9/2008	Documentation and dissemination: Development of instructional materials and release of software. Set up a forum to help disseminate information.

Section C – Project Details

1. Background

1.1. General background leading to this project. Please provide supporting documents, such as published papers by members of the project team or other researchers.

CONTEMPORARY COMPUTER ARCHITECTURE

Single and Multi-core Processors

In recent years, the design of microprocessors has undergone a remarkable transformation. Up until mid-2005, they were designed with minimal power consumption concerns and the industry was able to double the clock rate (and hence almost double performance) every 1.5 years. Ultimately, a power consumption issue arose, and continued progress along that trajectory could not be maintained.

Multi-core architectures were introduced to offer the potential for improved performance, and it is likely that the computer industry will double the number of cores on a chip every 1.5 years. Power consumption remains an issue and performance per watt is a better metric to reflect present concerns. Current dual-core processors, such as the 3 GHz Intel Core 2 Duo, have a maximum single precision floating performance of 24 Gflops per core (48 Gflops on a dual-core) [5].

Single core microprocessors are wasteful in terms of power and performance because the process of fetching, decoding and executing instructions along with caching data in a microprocessor is inherently sequential and the proportion of transistors that are actually used in the computation is very low. Architectures that support higher degrees of parallelism, such as multi-core processors, graphics processing units (GPUs) and field programmable gate arrays (FPGAs) can perform more floating point operations per second and hence offer the promise of improved performance with lower power.

FPGAs

A field programmable gate array (FPGA) is an array of logic gates in which the functionality and interconnection can be configured by downloading a bitstream into its memory. FPGA technology has been successfully applied to accelerate a large number of applications in areas as diverse as signal processing, communications, networking and robotics. The application of FPGA technology to computational problems is known as *reconfigurable computing*. FPGAs are becoming increasingly popular in high performance computing as they offer a method to design specialized processing units for a given application. They are among the fastest-growing electronic devices in recent years, and are used in an increasing range of applications as their speed, capacity and capabilities improve. They are beginning to appear as accelerators in supercomputers such as the Cray XD1, SRC-7 and SGI's RASC (Reconfigurable Application Specific Computing) unit. Furthermore, the FPGA and embedded systems markets are large and growing. Xilinx announced in April 2005 that it had achieved US$3B in sales of its Virtex device and, according to Business Communications Company Inc., embedded systems constitute a US$45.9B dollar market in 2004 which will increase to $88B by 2009. Recent work on dot product, matrix-vector and matrix multiplication [1] concluded that FPGAs will soon be able to significantly outperform modern microprocessors because of advantages in memory bandwidth and floating point performance.

One aspect of our own research focuses on the application of reconfigurable computing to domain specific problems. This involves developing improved computer architectures and arithmetic systems to improve performance over conventional techniques for a large number of different applications in cryptography, signal processing, scientific computing and financial engineering (refer to http://www.cse.cuhk.edu.hk/~phwl/publications.html for a complete list of papers). As examples, we developed FPGA-based architectures and methodologies that enabled a performance improvement of 4.5 times over a Pentium 4 processor that for a single precision implementation of the LINPACK benchmark [2] and for an interest rate derivative, a 25 times performance improvement over a Pentium was achieved [3]. These were the first reported implementations of LINPACK and derivative pricing applications. An ACM/IEEE 2006 Supercomputing conference Gordon Bell Prize finalist used FPGAs to act as accelerators for the N-body problem. The system was able to sustain 15.39 GFLOPS and its price/performance ratio was a remarkable $158/Gflops [7].

GPUs
Graphics processing units are ubiquitous in desktop platforms and serve to offload graphics processing requirements from the main processor. As an example, the Nvidia 7900 GTX has 24 execution units each capable of performing two single precision multiply-add operations on 4-element vectors at 650 MHz, giving a maximum performance of 250 Gflops. Their architectures support single precision floating point data types and recent results show that using an Nvidia 7900 GTX GPU, a 2-5 times speed improvement over a Pentium 4 processor can be achieved for sorting, fast Fourier transform and matrix multiplication [4].

PROGRAMMING ISSUES
Although the three technologies (multi-core, FPGA and GPU) have extremely impressive peak performance figures, unfortunately our ability to program them is very limited. Compilation of a standard serial C implementation of a program results in performance typically less than one tenth of the peak value. For example, Clustertech's optimized pricing of an American interest rate option achieves 0.27 Gflops on a Pentium 4 with peak performance of 2.3 Gflops. Utilizing the hardware effectively requires intimate knowledge of its architecture, means to parallelize the application and rethinking of the algorithms and optimization schemes. The general solution to this problem, i.e. a parallelizing compiler that can take a sequential program and produce efficient code for these architectures, has not been solved despite many decades of research.

Instead, domain-specific languages and libraries have been developed for these architectures. They include OpenMP and MPI for parallel microprocessors; Impulse C and Handel-C for FPGAs; and CUDA and CTM for GPUs. Companies such as RapidMind and PeakStream offer C++ libraries which use templates to generate code for multiple targets including multi-cores, GPUs and the Sony/IBM Cell processor. IEEE Spectrum has named the RapidMind Development Platform one of the 5 best technologies of the year [6]. These lower-level tools expose the execution units and memory system to the programmer so they can be used efficiently. Their main drawback is that they require specialized knowledge of the underlying architectures for effective use. Our approach to this problem is to limit our work to a specialized but important domain, and to hide these complications by placing this knowledge in the compiler plus automatic tuning system.

INVESTIGATOR'S BIOGRAPHIES
Prof. Philip Leong received the B.Sc., B.E. and Ph.D. degrees from the University of Sydney in 1986, 1988 and 1993 respectively. In 1993 he was a consultant to ST Microelectronics in Milan, Italy working on advanced flash memory-based integrated circuit design. From 1994-1997, he was a lecturer at the University of Sydney. Since 1997, he has been with the Department of Computer Science and Engineering at the Chinese University of Hong Kong where he is a Professor and the director of the Custom Computing Laboratory. Professor Leong is also Visiting Professor at Imperial College, London. He was program chair of the FPT and FPL conferences and an associate editor for the ACM Transactions on Reconfigurable Technology and Systems. The author of more than 80 technical papers and 4 patents as well as the recipient of the 2006 FPT conference best paper award, his research interests include reconfigurable computing, signal processing, computer architecture, computer arithmetic and biologically inspired computing.

Prof. Jimmy Lee graduated on the Dean's Honours List with an Honours BMath degree, majoring both in Applied Mathematics and Computer Science, from the University of Waterloo, Canada, in 1987, and obtained his MMath degree in Computer Science in 1988 also from Waterloo. He then moved to the University of Victoria, Canada, and completed his doctoral study in 1992. During his graduate studies, he was involved in research projects funded by IBM (Canada) and the Canadian Institute of Robotics and Intelligent Systems. Upon graduation, he joined The Chinese University of Hong Kong, where he is now a Professor in the Department of Computer Science and Engineering. His research focuses on the theory and practice of constraint satisfaction and optimization with applications in combinatorial optimization, scheduling, and resource allocation. Author of over 80 technical papers and an Associate Editor of the Journal of Discrete Algorithms, he is particularly interested in problem modeling, stochastic local search, and over-constrained problems. Starting 2005, he is also the Associate Director (Research and Development) at the Centre for the Advancement of Information Technology in Education (CAITE) under the umbrella of the Hong Kong Institute of Educational Research.

Prof. Wayne Luk received his M.A., M.Sc. and D.Phil. in engineering and computing science from Oxford University. Currently Professor of Computer Engineering at Imperial College London, he founded and leads the Custom Computing Group there. He is a Visiting Professor at Stanford University and Queen's University Belfast. He is a founding Co-Editor-In-Chief of ACM Transactions on Reconfigurable Technology and Systems, and has served on the Program Committee of many international conferences such as FPL, FCCM, FPGA, FPT, DATE, ASP-DAC and ASAP, and was chair on six occasions. He has been an author or editor for 6 books, 3 special journal issues, a patent and published over 120 research papers in peer reviewed journals and conference proceedings. His research interests include reconfigurable computing and high level synthesis.

Dr Kenneth Chow is a founder and the Chief Operating Officer at Cluster Technology. He obtained his B.Sc and Ph.D from Caltech and Cornell respectively. He has over 12 years of experience in financial modeling and risk management. Before founding Cluster Technology, he led the quantitative research group at Goldman Sachs Asia.

Chris Chi-Chiu Cheung received a Bachelor of Science with first class honours in physics and an MPhil in computer science and engineering from the Chinese University of Hong Kong in 1995 and 1997 respectively. He is currently a senior project manager at Cluster Technology Ltd. His research interests include parallel computing, distributed computing, and software engineering.

1. Keith D. Underwood, K. Scott Hemmert: Closing the Gap: CPU and FPGA Trends in Sustainable Floating-Point BLAS Performance. Proceedings FCCM, pp. 219-228, 2004
2. Kieron Turkington, Konstantinos Masselos, George A. Constantinides and **Philip Leong**, FPGA based Acceleration of the LINPACK Benchmark: A High Level Code Transformation Approach, *Proceedings of the Eleventh International Workshop on Field Programmable Logic and Applications (FPL'06)*, Madrid, pp. 275-380, 2006
3. G. L. Zhang, **P. H. W. Leong**, C. H. Ho, K. H. Tsoi, **C. C. C. Cheung**, D-U. Lee, R. C. C. Cheung and **W. Luk**, Reconfigurable Acceleration for Monte Carlo based Financial Simulation, *Proceedings of the International Conference on Field Programmable Technology (FPT),* Singapore, pp. 215-222, 2005
4. Govindaraju, N. K., Larsen, S., Gray, J., and Manocha, D. 2006. A memory model for scientific algorithms on graphics processors. In *Proceedings of the 2006 ACM/IEEE Conference on Supercomputing* (Tampa, Florida, November 11 - 17, 2006). SC '06.
5. http://download.intel.com/technology/architecture/sma.pdf
6. http://www.spectrum.ieee.org/print/4837
7. Kawai, A. and Fukushige, T. 2006. Gordon Bell finalists II-$158/GFLOPS astrophysical N-body simulation with reconfigurable add-in card and hierarchical tree algorithm. In *Proceedings of the 2006 ACM/IEEE Conference on Supercomputing* (Tampa, Florida, November 11 - 17, 2006). SC '06.

1.2. Any previously related project(s) undertaken by your organization and/or key project team members in the past five years and supported by the ITF or any other funding sources in Hong Kong and around the world? If yes, please briefly describe the relevant/related project(s) and the source(s) of funding obtained for it (them).

RGC Earmarked Grant, "(CUHK 4333/02E) Floating Point Computation using FPGA Clusters", 2002-2004, HK$ 413,404, *Philip Leong (PI) and Wayne Luk (co-I)*

We developed a tool which automates the translation of a high level algorithmic description to an FPGA implementation. In order to minimise the hardware resources, computations are done using floating point or fixed point formats with the minimum necessary precision. The system consists of an arbitrary precision floating point class for simulation; a profiler used to determine the minimum precision required for every signal in the hardware implementation; an optimiser which finds the minimum number of bits for a specified degree of accuracy; and a compiler which takes the information collected by the system and outputs synthesisable VHDL code. The resulting designs have hardware requirements which exceed that of a single FPGA chip. To address this problem, a cluster of 8 Linux computers, each equipped with a high capacity FPGA accelerator board was utilised. The results were useful for developing future FPGA clusters with hundreds of millions of system gates and beyond. The ultimate goal of this research was to pioneer the application of FPGA based coprocessors to large scale floating point intensive applications. As a first application, a Monte Carlo simulation accelerator for the pricing of financial derivatives was developed. We also developed a number of high performance Gaussian random number generators in this work.

1. G. L. Zhang, **P. H. W. Leong**, C. H. Ho, K. H. Tsoi, **C. C. C. Cheung**, D-U. Lee, R. C. C. Cheung and W. Luk, Reconfigurable Acceleration for Monte Carlo based Financial Simulation, *Proceedings of the International Conference on Field Programmable Technology (FPT)*, Singapore, pp. 215-222, 2005

ITF Matching Grant UIM/68, "Interface software for Cluster Computing," 2001-2004 $4,371,000 (total $10,512,000). Philip Leong (PI) and Cluster Technology (Industrial Sponsor).

We developed an industrial quality library to facilitate the parallelization of financial engineering applications and is described in detail in [1]. High performance template-based C++ libraries for the solution of partial differential equations, Monte Carlo simulation and artificial neural networks were developed. Linear scaling was achieved on clusters up to 64-nodes. This is now a commercial software product being marketed by Cluster Technology (www.clustertech.com). Licenses have been sold to HSBC and another major Asian bank. Furthermore, the expertise developed through this grant has led Clustertech to become a leader in high performance computing consulting business and helped to attract high profile clients such as the Environmental Protection Department of Hong Kong, the Hong Kong Observatory, the Hong Kong Jockey Club, Pan Asia, HSBC and ITV (UK). For more information, see http://clustertech.com/index.php?Products:CPE.

1. M.P. Leong, **C.C. Cheung**, C.W. Cheung, P.P.M. Wan, K.H. Leung, W.M.M. Yeung, W.S. Yuen, **K.S.K. Chow**, K.S. Leung and **P.H.W. Leong**, A Parallel Library for Financial Engineering Applications, *IEEE Computer*, Vol. 38, No. 10, pp. 70-77, 2005

HK Innovation and Technology Fund SERAP Grant S/P779/04B -- "A Software Platform to Enable the Commercial Deployment of High Performance Computing Clusters", 2005—2006, HK$1,606,300 Ken Chow, Wai-Mo Suen and Philip Leong)

This proposal involved the development of the Clustertech HPC Environment Software Stack (CHESS) which implements the execution environment for Linux-based HPC clusters. CHESS contains a rich set of components covering all your HPC cluster needs, including parallel middleware, job scheduling, resource management, real-time system monitoring and development tools. These components are tightly integrated with the CHESS Information Infrastructure (CHESS-I2), an infrastructure enabling a centralized management of the HPC cluster. Clustertech is partnering with Dell Computers to sell CHESS-equipped clusters in China and there are currently 31 CHESS clients: Anhui University, Beijing City University, Beijing Normal University, Beijing University of Science and Technology, Beijing University of Technology, China Aviation Industry Corporation (606 Suo), China University of Mining and Technology, Automation, China University of Petroleum, China University of Petroleum UPC, CNOOC, Dalian Maritime University, Dalian University of Technology DLUT, Dalian University, Chemistry, GeoTomo China GeoTomo, Huazhong Normal University, Institute of Geology and Geophysics, Institute of Microbiology, CAS, Institute of Plasma Physics, CAS, National Marine Data & Information Service, Ocean University, Petrochina Northwest, Shanghai Institutes for Biological Sciences, Shanxi University, Physics, Sichuan University Chemistry, Sinopec Hunan, Southwest Normal University Chemistry, Thomson Broadband R&D

(Beijing) Co. Ltd, Tsinghua University, University of Science and Technology of China, Xiamen University and Xiamen University Chemistry. For more information, see http://clustertech.com/index.php?Products:CHESS.

UK Engineering and Physical Sciences Research Council, EPSRC (EP/D060567/1)
"Reconfigurable Architectures for Floating Point Applications", 2006-2008, HK$7,212,000, Philip Leong (PI) and Wayne Luk (co-I)

This proposal involves the VLSI design of field programmable gate array (FPGA) architectures optimised for floating point applications. A modelling environment will be developed to explore tradeoffs associated with high level architectural decisions, and floating point benchmark problems will be studied. Software which can take a high level description of a floating point FPGA (FPFPGA) and produce an application specific integrated circuit (ASIC) will also be developed. This software will be used to create both general and domain specific integrated circuits.

Hong Kong Research Grants Council (CUHK 413105)
"Breaking Symmetries of Indistinguishable Values in Constraint Satisfaction", 2005-2007 HK$692,480, Jimmy Lee (PI)

Constraint satisfaction problems (CSPs) sometimes contain both variable symmetries and value symmetries, causing adverse effects on CSP solvers based on tree search. As a remedy, symmetry breaking constraints are commonly used. While variable symmetry breaking constraints can be expressed easily and propagated efficiently using lexicographic ordering, value symmetry breaking constraints are often difficult to formulate. We propose two methods of using symmetry breaking constraints to tackle value symmetries. First, we show theoretically when value symmetries in one CSP correspond to variable symmetries in another CSP of the same problem. We also show when variable symmetry breaking constraints in the two CSPs, combined using channeling constraints, are consistent. Such results allow us to tackle value symmetries efficiently using additional CSP variables and channeling constraints. Second, we introduce value precedence, a notion which can be used to break a common class of value symmetries, namely symmetries of indistinguishable values. While value precedence can be expressed using inefficient if-then constraints in existing CSP solvers, we propose efficient propagation algorithms for implementing global value precedence constraints. We also characterize several theoretical properties of the value precedence constraints. Extensive experiments are conducted to verify the feasibility and efficiency of the two proposals.

1. Y.C. Law and **J.H.M. Lee**. Symmetry Breaking Constraints for Value Symmetries in Constraint Satisfaction. (to appear) *Constraints*, 2007

The Macao Water Supply Co. Ltd. Research Contract.
"Study for Optimization of Operation during the Salinity Period", 2004-2006, HK$191,100.00, Jimmy Lee (PI)

The project is to help the water company of Macau to control the salinity level of potable water. During November of each year to March of the next year, the level of the source of water supply to Macau is low causing high salinity level in potable water. Care must then be taken to control the opening of pumps at the water source and the transfer of water between reservoirs to control the salinity level of potable water for the general public. The goal of the project is to construct an optimization engine to help optimize the operations at the various pumping station and reservoirs so as to maximize the number of days that the potable salinity level is below a user acceptable level using a constraint programming approach.

1. C.W. Choi and **J.H.M. Lee**. Controlling Salinity in a Potable Water Supply System Using a Constraint Programming Approach, *Proceedings of the 17th IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2005)*, pages 104-108, November, Hong Kong, 2005

Hong Kong Research Grants Council (CUHK 421904)

"Utilizing Useful Redundancy Information to Increase Search Efficiency in Constraint Satisfaction", 2004-2006
HK$339,414, Jimmy Lee (PI)

Combining mutually redundant models with channelling constraints increases constraint propagation. However, the extra computation efforts of the additional variables and constraints may outweigh the gain of reduction in search space. In fact, many of the constraints in redundant modelling are not only logically redundant but also propagation redundant and hence cannot further reduce search space. We give general theorems for proving propagation redundancy of one constraint with respect to channelling constraints and constraints in the other model for the domain consistency propagator. We define a broad form of channelling constraints that are covered by our approach. We illustrate, using problems from CSPLib (http://www.csplib.org/), how detecting and removing propagation redundant constraints can significantly speed up solving behaviour. We extend the results and cover also the case for set bounds consistency propagator.

1. C.W. Choi, **J.H.M. Lee**, and P.J. Stuckey. Removing Propagation Redundant Constraints in Redundant Modeling. (to appear) *ACM Transactions on Computational Logic*, 2007

1.3. Any other R&D work or projects similar to this application done or being carried out by other parties in Hong Kong and around the world? If yes, please set out the findings and explain how your approach is compared to others in terms of technological superiority, production costs, market acceptability, etc.

Although previous work on parallelizing compilers, automatic tuning, using low precision for speed and high precision for checking, reconfigurable computing, GPU programming and multi-core programming are all active research areas (and reviewed in section 1.1), we are not aware of domain-specific parallelizing compilers of the type that we are proposing.

1.4. Any pilot work has already done by your organization and/or project team members in preparation for this project? If yes, please describe the work done.

None. However we believe our experience in related fields such as MPI programming, reconfigurable computing, computer architecture, performance optimization of software programs, cluster computing, optimization and derivative pricing puts us in a unique position for this project.

1.5. Any request for funding support for this application previously rejected by ITF, Research Grants Council (RGC) and/or other Government-administered funding schemes? If yes, please set out the project reference of the previous application.

N/A

1.6. If this application is a previously rejected application under ITF, any major differences of this application vis-à-vis the previous one? Please explain how these differences have addressed the previously raised concerns.

N/A

1.7. Any attempt made or being made to seek funding support for this project from other funding sources in Hong Kong and around the world (e.g. RGC, Applied Research Fund, Environment and Conservation Fund, etc.)? Please advise any known results of the application(s).

None.

2. Implementation Approach

2.1. Please elaborate on the technology to be developed and/or the innovative use of existing technologies. The brief information provided in Section A is relevant.

Most developers of Monte Carlo Simulations (MCS) and optimization, including banks, insurance companies and scientific organizations, do not have backgrounds in computer architecture or parallel programming. Thus the barrier for entry to utilize array technologies is very high. The aim of this work is to develop tools that enable them to do so, while maintaining a familiar, serial programming model.

The issues that need to be addressed in our approach are:

- High level development – the source language is standard C.
- Efficiency – detailed analyses of execution unit utilization, dependencies, pipelining, memory hierarchies etc need to be made in order to ensure that the computation is efficiently mapped to the hardware.
- Accuracy – although most software uses double precision floating point, single precision computations are normally faster as they require lower memory bandwidth and map better to array technologies.
- Portability – the same code should be able to run on multiple targets with the highest performance and accuracy.

Most modern high performance software is written in C or C++ and we believe a successful product must allow for easy porting of existing applications. In order to achieve **portability**, we will develop an optimizing C-to-C (otherwise known as a source-to-source) compiler that targets our application domain. We have experience in developing this type of tool for reconfigurable computing applications [9,10]. The user will write a serial description of code to be executed on an array technology, and our library will generate efficient code for the desired platform (multi-core, GPU or FPGA). Different programming languages will be used as the output language of our compiler to support the different array technologies. As the same C-based **high level** program can be compiled to different targets, it will be easy to generate and execute a program on multiple array technologies and choose the most efficient implementation.

To achieve **efficiency**, fast primitive libraries must be used. For MCS, the Gaussian random number generators are extremely important. We are research leaders in this field [1-4] and will develop optimized primitives for the different array technologies. We have shown that MCS applications can be parallelized by pipelining [5] or by executing different paths on different processors [6] and believe that our selected optimization problems are similar. Furthermore, many optimizations are processor-specific and tuning of optimization parameters difficult. We will employ an automatic generation and parameter tuning scheme, such as is used in the widely available ATLAS (Automatically Tuned Linear Algebra) package [7] to address this problem. The resulting system will be the first MCS library that can target different array technology hardware. We also believe that it will achieve efficiency surpassing all previously reported systems.

Most computations for financial and risk management applications do not require double precision **accuracy**. However, in practical applications, some guarantee of accuracy is mandatory. In order to achieve the highest possible efficiency, we hence propose to do most computations in single precision floating point. Our library will contain verification code that uses double precision to verify the answer by computing a limited number of paths using double precision and comparing the statistics of these with those in single precision. To the best of our knowledge, we will be the first to use this technique for MCS. The use of single precision operations to achieve double precision results is an important recent research area [8] and we hope to use recent results in this area in our work as well as contribute new algorithms of this type.

To summarize, by combining the novel ideas above, we will develop a compiler that is capable of generating optimized code for different array technology platforms with little human intervention. Such a system will be the first multi-target, self-testing, self-tuning system of its kind. Although this project will focus on MCS and optimization problems, we believe this paradigm can be extended to almost any other problem in high performance computing.

1. David Thomas, Philip Leong, Wayne Luk and John D. Villasenor, A Survey of Gaussian Random Number Generators, accepted for publication in ACM Computing Surveys, March 2007 (preprint available from http://www.cse.cuhk.edu.hk/~phwl/mt/public/archives/papers/grng_acmcs07.pdf)
2. D. Lee, J.D. Villasenor, W. Luk and P.H.W. Leong, A hardware Gaussian noise generator using the Box-Muller method and its error analysis, IEEE Transactions on Computers, volume 55, number 6, pages 659-671, Jun 2006
3. Dong-U Lee, Wayne Luk, John D. Villasenor, Guanglie Zhang and Philip H.W. Leong, A Hardware Gaussian Noise Generator using the Wallace Method, IEEE Transactions on VLSI Systems, Vol. 13, No. 8, pp. 911-920, August 2005
4. Philip H.W. Leong, Ganglie Zhang, Dong-U Lee, Wayne Luk and John D. Villasenor, A comment on the Implementation of the Ziggurat Method, Journal of Statistical Software, Vol. 12, Issue. 7, 2005 (http://www.jstatsoft.org)
5. G. L. Zhang, P. H. W. Leong, C. H. Ho, K. H. Tsoi, C. C. C. Cheung, D-U. Lee, R. C. C. Cheung and W. Luk, Reconfigurable Acceleration for Monte Carlo based Financial Simulation, Proceedings of the International Conference on Field Programmable Technology (FPT), Singapore, pp. 215-222, 2005
6. M.P. Leong, C.C. Cheung, C.W. Cheung, P.P.M. Wan, K.H. Leung, W.M.M. Yeung, W.S. Yuen, K.S.K. Chow, K.S. Leung and P.H.W. Leong, A Parallel Library for Financial Engineering Applications, IEEE Computer, Vol. 38, No. 10, pp. 70-77, 2005
7. http://www.netlib.org/atlas/index.html
8. Langou, J., Langou, J., Luszczek, P., Kurzak, J., Buttari, A., and Dongarra, J. 2006. Tools and techniques for performance---Exploiting the performance of 32 bit floating point arithmetic in obtaining 64 bit accuracy (revisiting iterative refinement for linear systems). In Proceedings of the 2006 ACM/IEEE Conference on Supercomputing (Tampa, Florida, November 11 - 17, 2006). SC '06.
9. T. Todman, J.G. de F. Coutinho and W. Luk, Customisable hardware compilation, The Journal of Supercomputing, vol. 32, no. 2, pp. 119-137, May 2005
10. C.H. Ho, P.H.W. Leong, K.H. Lee, K.H. Tsoi, R. Ludewig, P. Zipf, A.G. Ortiz and M. Glesner Fly - A Modifiable Hardware Compiler, Proceedings of the Nineth International Workshop on Field Programmable Logic and Applications (FPL'02), Montpellier, LNCS vol. 2438, pp. 381-390, 2002

2.2. Please elaborate with technical details on each project milestone. The brief information provided in Section B is relevant.

Milestone 1: Hardware evaluation.
Upon commencement of the project, we will study GPU cards from the main vendors, Nvidia and ATI and select the most suitable given our budget. For the reconfigurable computing platform, we will study recent low latency, front-side bus based platforms such as those made by DRC (http://www.drccomputer.com) and Xtreme Data (http://www.xtremedatainc.com/) as well as more traditional ones such as being made by Celoxica (www.celoxica.com). Equally important is availability and quality of software development kits which will be taken into account in a purchasing decision. Orders for the hardware will be placed.

Milestone 2: Libraries and Benchmarks
In order to produce efficient implementations, efficient libraries are required. We will also develop hand-optimized implementations of primitives such as random number generators for the selected platforms. These will consider low level issues such as degree of parallelism, memory hierarchy and data dependencies. Our study is limited to Monte Carlo simulation and optimization applications. We will develop a set of benchmarks for Monte Carlo simulation, simulated annealing, genetic algorithm and grid search.

Milestone 3: Compiler
We will first produce hand-optimized implementations of the benchmark programs and ensure that they achieve good performance on the multi-core, reconfigurable computing and GPU platforms. We will then develop a simple compiler that can take the kernels of the miletone 2 benchmark applications and generate optimized code for the different array technologies. It is likely that existing infrastructure such as LLVM (http://llvm.org/) which can produce C output will be used. LLVM uses the GCC front-end, ensuring portability and standards compliance. It supports most modern

optimizations including inlining, IP constant propagation, dead argument elimination, dead global elimination, making global variables constants, by-ref to by-value argument promotion, etc. Its license allows us to sell our resulting modified software.

This milestone is the main milestone for the project. The issues to be addressed are:

- How to allocate memory. Ideas from Stanford's Sequoia [1] will be used.
- How to achieve self-tuning of the code. This is relatively easy as compiler decisions such as degree of loop unrolling, loop blocking, execution unit assignment etc can simply be parameterized and treated as an optimization problem.
- How to check accuracy. For Monte Carlo simulations, we can compute a limited number of paths in double precision and check that they are converging to the same single precision values. For optimization, we can check that evaluation of the cost function in both single and double precision is consistent.

Milestone 4: Documentation and marketing
Full documentation is beyond the scope of a project given its short duration. Instead, we will produce conference publications to disseminate the results of our work. Material will also be developed for our website along with a tutorial and sample applications. We will approach existing Clustertech clients to be alpha testers of this work.

1. http://www.stanford.edu/group/sequoia/cgi-bin/

3. Target Results and Benefits

3.1. Who the target user groups are?

In this, the first stage of the project, our target users are banks, insurance companies, scientific organizations and researchers involved in the development of high performance software for array technologies. Furthermore, to focus on the Chinese market, our documentation and tutorials will be available both in English and Chinese. This, together with Clustertech's base of close to 50 clients in the greater China region, will allow us to gain a solid base of users in the first two years.

Ultimately, we hope to have a compiler that would be valuable to all software developers. As multi-cores and GPUs are already dominant and reconfigurable computing is rising in acceptance for high performance computing applications, tools of this kind will become necessary for all software developers concerned with performance.

3.2. The estimated number of companies/customers to adopt/use the new facilities/ technology/product/service/ process.

In the Monte Carlo simulation and optimization applications, we hope to find at least 5 existing Clustertech clients to use the developed compiler. Through our free download, we hope to have thousands of users, some of which will become commercial clients. In several years, when more domains are covered, we hope to expand both of these estimates by a factor of 10.

3.3. Expected number of beneficiaries vis-a-vis the population of the target user group(s).

We aim for thousands of people to benefit directly from this work, and much higher numbers in the future when we expand to other domains.

3.4. Details of the dissemination plan, including the activities that would be undertaken to promote the project deliverables to the target users and the justifications for proposing and organizing such activities. The brief information provided in Section A is relevant.

1. We will publish papers in peer reviewed conferences and journals to make others aware of our work in this area.
2. We will implement a project website with free download of the software for non-commercial use.
3. Clustertech will create and maintain a forum (Chinese and English) on which the public can make enquires and ask questions regarding the software. The Chinese University and Imperial College will participate on the forum.

3.5. The proposal's contribution to Hong Kong’s innovation and technology development.

Multi-cores, reconfigurable computing and GPUs are a disruptive technology that could change the balance of high performance computing just as powerful microprocessors took the market from vector processors in the past. We believe that a great opportunity exists for Hong Kong to take a lead in this area, and this project can serve not only to develop a useful tool, but also to train students in this area. As our compiler lowers the knowledge needed to utilize array technologies, it will enable software developers in Hong Kong and mainland China to harness the available power of array technologies and use them to develop innovative software. Our ability to provide documentation in Chinese and English as well as our connections to China gives us an advantage over US and European companies. Application areas that come immediately to mind are financial engineering, risk management, enterprise resource management, game programming and pollution and weather modeling. This work is applicable to any companies that have a need for improved performance in their applications.

The Hong Kong Monetary Authority (HKMA) has recently placed new risk management reporting requirements on banks according to the "advanced approaches promulgated under Basel II" [1]. In order to stress test a bank's market and credit risk, there will be new requirements that require large scale Monte Carlo Simulations [2]. Monte Carlo Simulations will be increasingly used as a result of this new ordinance.

1. http://www.info.gov.hk/hkma/eng/viewpt/20050714e.htm
2. http://www.info.gov.hk/hkma/eng/research/RM15-2006.pdf

3.6. Others (please specify).

N/A

4. Collaborations with Other Organisations

4.1. Any collaboration with other organizations? If so, please elaborate on the form of such collaborations.

Clustertech will provide both cash and in-kind contributions to this project. A letter of support from Clustertech is included in this proposal.

Clustertech and Dell Computers are collaborating on building up the high performance computing sector in mainland China. A total of 31 cluster installations have been made to date.

Imperial College and Clustertech are collaborating with the Chinese University of Hong Kong on research in the areas of financial engineering, reconfigurable computing and VLSI design. We have a track record of grants and publications in these areas.

4.2. Any special arrangements arising from such collaboration, e.g. licensing of intellectual property rights? If so, please elaborate?

Intellectual property will be owned by the Chinese University of Hong Kong. It will be licensed (non-exclusively) to Cluster Technology and any other interested parties for commercialization.

Profit sharing between the Chinese University and Cluster Technology is being negotiated.

5. Business Plan

(Please elaborate on the business plan (if any) for the project. It should address, without limitations, how to generate income after project completion; how to maintain or update the deliverables; and whether to develop spin-off units after the completion of the project.)

We intend that the Chinese University of Hong Kong and Imperial College continue research in this area after completion of the project. Maintenance, support and future development will be carried out by Clustertech.

Clustertech will potentially hire the project's staff upon completion of the project. It will approach its existing clients to obtain development and consulting work on array technologies. These clients total 50 in number and include the Environmental Protection Agency of Hong Kong, the Hong Kong Observatory, the Hong Kong Jockey Club, HSBC, ITV, a major Asian bank, Chinalink etc. We will also jointly find new clients via our free software and Clustertech's marketing arm.

6. Patent Plan

(Please complete the following table if you have any intention to patent any of the process(es) or product(s) to be developed under the project.)

Name of Patentable Item	Country where Registration will be filed
N/A	

7. Other Information in Support of the Application

N/A

Section D – Applicant Organization and Collaborating Parties

(When an application is approved, the information in this section will be published at the ITF website (denoted with #) and/or carried forward as pre-filled data in the related progress/final reports.)

1. Applicant Information#

Name in English	:	The Chinese University of Hong Kong
Name in Chinese	:	香港中文大學
Year of Establishment	:	
Nature of Business	:	Higher Education Institution
Registered Address	:	Shatin, New Territories, Hong Kong
Telephone Number	:	852-26961836
Fax Number	:	852-26037414
Email Address	:	agneslam@cuhk.edu.hk
Webpage	:	www.cuhk.edu.hk
Contact Person - Name	:	Miss Agnes Fung-yee LAM
Contact Person - Position	:	Executive Officer, RAO

2. Collaborating Parties#

No	English Name (Chinese Name)	Role in the Project	Address / Webpage (if any)	Contact Person	Tel No / Fax No / Email
1	Cluster Technology Limited (聯通計算科技有限公司)	Sponsoring Organization	Suites 4316-19, 43/F, Jardine House, 1 Connaught Place, Central, Hong Kong/ www.clustertech.com	Dr Kenneth CHOW	852-90291975/ 852-29942101/ chowskei@clustertech.com
2	Imperial College (-)	Collaborating Organization	180 Queen's Gate London/ www.ic.ac.uk	Prof Wayne Luk	+44 207 594 8313/ +44 207 581 5024/ wl@doc.ic.ac.uk

Section E – Project Team

(When an application is approved, the information in this section will be published at the ITF website (denoted with #) and/or carried forward as pre-filled data in the related progress/final reports.)

1. Project Coordinator#

Project Role	:	Project coordinator + Principal investigator
☒**Key member** ☒**CV included** ☐**To be Paid by the Project**		
Name in English	:	Dr Philip Heng-wai LEONG
Name in Chinese	:	梁桓惠博士
Position	:	Professor
Department (if any)	:	Department of Computer Science and Engineering
Organization Name	:	The Chinese University of Hong Kong
Telephone Number	:	852-26098414
Fax Number	:	852-26035024
Email Address	:	phwl@cse.cuhk.edu.hk
Main Task	:	Responsible for the overall direction and technical details of the project. Supervise all members and liaise with Cluster Technology and Imperial College.
Organization Webpage	:	www.cuhk.edu.hk

2. Other Team Members#

No	Key Member?	Name (Chinese name) Role in the Project	Main Task	Position or Project Post/Rank Department Organisation	With CV?
1	*	Prof Jimmy Ho Man Lee (李 浩 文) Investigator	Will work with Prof Leong on management and technical directions of the project. In particular, he will contribute his expertise in optimization and programming languages.	Professor Computer Science and Engineering The Chinese University of Hong Kong	Yes
2	*	Mr Kenneth CHOW (周信基) Deputy project coordinator	Advise on project direction and marketing of results.	Chief Operating Officer - Cluster Technology Limited	Yes
3	*	Prof Wayne Luk (-) Investigator	Advise on high level synthesis, computer architecture and applications	Professor of Computer Engineering Computing Imperial College	Yes
4	*	Mr Chris Cheung (張志超) Investigator	Contribute advice on software engineering aspects of the project as well as derivative pricing and optimization applications.	Senior Project Manager - Cluster Technology Limited	Yes

Section F – Budget for the Project

(Please provide full justifications for each sub-item under the budget items "Manpower", "Equipment" and "Other Direct Costs". The rationale behind any projected income or expenditure has also to be given. In case certain goods or services are intended to be procured from one company/organization/individual, please provide the details, relationship between the applicant(s) and the company/organization/individual (if any) and justifications for not following the open procurement procedures set out in the "Guide to the Innovation and Technology Fund". In addition, quotation/price information should be provided in justifying expenditure items of $500,000 or more.)

1. Expenditure

(Please ensure that all expenditure items must be incurred between the commencement and completion dates of the project.)

1.1. Manpower

Key member	Post/Rank	No. of staff (A)	Duration (B) (man-month)	Monthly rate or equivalent (C) ($'000)	Total (A)*(B)*(C) ($'000)	Justification
	Research Assistant (MPSB11)	2.0	12.0	14.380	345.120	[1]
	Senior Research Assistant (MPSB19)	1.0	12.0	21.725	260.700	[2]
				Sub-total (I):	**605.820**	

1.2. Equipment

(For each sub-item under "Equipment", apart from providing justifications for its procurement, please also state in the explanatory notes whether similar equipment is available for sharing within the applicant organization or with other ITF recipient organizations, and if so, the reason why the existing equipment cannot be used for this project.)

Key equipment	Item	Quantity (A)	Unit cost (B) ($'000)	Total (A)*(B) ($'000)	Justification
*	Graphics processing unit	1.0	20.000	20.000	[3]
*	Reconfigurable computing card	1.0	100.000	100.000	[4]
*	Host computers	4.0	10.000	40.000	[5]
*	Software	1.0	10.000	10.000	[6]
			Sub-total (II):	**170.000**	

1.3. Other Direct Costs

(In case external consultants are required for the project, please set out clearly in the explanatory note the justifications for engaging the consultants and the expected time commitment of the consultants under the project.)

Item	Quantity (A)	Unit cost (B) ($'000)	Total (A)*(B) ($'000)	Justification
Audit fees	1.0	5.000	5.000	[7]
Conference	1.0	20.000	20.000	[8]
		Sub-total (III):	**25.000**	
		Total((I)+(II)+(III)):	**800.820**	

[1] Required to develop libraries on multi-core, GPU and FPGA platforms
[2] Responsible for high level design of MCS system
[3] Required for development purposes
[4] For development of FPGA-based accelerators. The price quoted is the market price for one unit.
[5] Three machines will be used to host the multicore, GPU and reconfigurable computing platforms. The other will be a centralized server for this project's software. As they will be used for benchmarking and development, no sharing opportunities exist.
[6] Purchase GPU programming and software libraries for programming multi-cores, FPGAs and GPUs.
[7] Auditing expense for project less than $1M in total.
[8] To disseminate results of our work

2. Amount of Sponsorship

(Please provide proof of the sponsorship in the form of annex to this application.)

Sponsoring Organisation	Cash	Equipment	Consumables	Total
		(in cash-equivalent)		
	(A) ($'000)	**(B)** ($'000)	**(C)** ($'000)	**(A)+(B)+(C)** ($'000)
Cluster Technology Limited	100.000	0.000	0.000	100.000
			Total:	**100.000**

3. Income

(Please set out all income, except sponsorship, in this section.)

Item		Total ($'000)
	Total:	**0.000**

4. Administrative Overheads

(Please budget not more than 15% of the net of total expenditure, sponsorship and income in this section if this is a proposal submitted by a local university.)

Item		Total ($'000)
University overhead		105.123
	Total:	**105.123**

5. Net Amount Requested from the Innovation and Technology Fund

Total Expenditure	Total Sponsorship	Total Income	ITF Funding (net of overhead)	ITF Funding (overhead)	Net Requested Amount from ITF
(A) ($'000)	**(B)** ($'000)	**(C)** ($'000)	**(D)=(A)-(B)-(C)** ($'000)	**(E)** ($'000)	**(D)+(E)** ($'000)
800.820	**100.000**	**0.000**	**700.820**	**105.123**	**805.943**

Section G – Classification of the Project

(When an application is approved, the information in this section will be published at the ITF website.)

1. Research Theme

Regular

2. Project Type

Product/Process/Technology Development/Improvement

3. Technology Area

3.1. Primary Area

Information Technology

3.2. Other Areas (if any)

3.2.1. Electrical and Electronics
3.2.2.
3.2.3.

4. Industrial Sector

4.1. Primary Sector

Information Technology

4.2. Other Sectors (if any)

4.2.1. Electrical and Electronics
4.2.2.
4.2.3.

5. Other Attributes (if any)

(Please feel free to choose more than one attribute.)

5.1. Environment-related ☐
5.2. Quality-related ☐
5.3. Particularly SME-related ☒

Section H – Attachments for the Project

Annex No.	Section No.	Paragraph No.	File Name
1	C		cv_leong.doc
2	C		cv_jlee.doc
3	C		cv_chow.doc
4	C		cv_luk.doc
5	C		cv_cheung.doc
6	C		clustertech_letterofsupport.pdf

I hereby declare that:

(a) all factual information provided in this application as well as the accompanying information reflects the status of affairs as at the date of submission. I shall inform the Secretariat of the Innovation and Technology Fund immediately if there are any subsequent changes to the above information; and

(b) the ideas of the proposed project are original without any constituted or potential act of infringement of the intellectual property rights of other individuals and/or organizations.

Authorized Signature with Organization Chop: ______________________________

Name of Signatory: ______________________________

Position: ______________________________

Name of Principal Applicant Organization: ***The Chinese University of Hong Kong***

Date (dd/mm/yyyy): ______________________________