



AARHUS  
UNIVERSITY

BUSINESS AND SOCIAL SCIENCES  
DEPARTMENT OF ECONOMICS AND BUSINESS



**CREATES**

Center for Research in Econometric Analysis of Time Series

# **Heterogeneous Computing in Economics: A Simplified Approach**

**Matt P. Dziubinski and Stefano Grassi**

**CREATES Research Paper 2012-15**

# HETEROGENEOUS COMPUTING IN ECONOMICS: A SIMPLIFIED APPROACH

MATT P. DZIUBINSKI AND STEFANO GRASSI

*CREATES, Department of Economics and Business, Aarhus University,  
Denmark*

**ABSTRACT.** This paper shows the potential of heterogeneous computing in solving dynamic equilibrium models in economics. We illustrate the power and simplicity of C++ Accelerated Massive Parallelism (C++ AMP) recently introduced by Microsoft. Starting from the same exercise as Aldrich et al. (2011) we document a speed gain together with a simplified programming style that naturally enables parallelization.

*JEL Classification:* C88.

*Keywords:* Code optimization; CUDA; C++; C++ AMP; Data parallelism; DSGE models; Econometrics; Heterogeneous computing; High-performance computing; Parallel computing.

## 1. INTRODUCTION

This paper shows the potential of heterogeneous computing in solving dynamic equilibrium models in economics. Heterogeneous computing refers to the use of different processing cores (types of computational units) to maximize performance.<sup>1</sup>

We rely on C++ Accelerated Massive Parallelism (C++ AMP), a new technology introduced by Microsoft, that allows to use accelerators, such as graphics processing units (GPUs), to speed up calculations. GPUs are a standard part of the current personal computers and are designed for data parallel problems, where we assign an individual data element to a separate

---

*E-mail address:* sgrassi@creates.au.dk.

*Date:* August 6, 2012.

We acknowledge financial support by the Center for Research in Econometric Analysis of Time Series, CREATES, funded by the Danish National Research Foundation.

<sup>1</sup>"A computational unit could be a general-purpose processor (GPP) – including but not limited to a multi-core central processing unit (CPU), a special-purpose processor (i.e. digital signal processor (DSP) or graphics processing unit (GPU), a co-processor, or custom acceleration logic (application-specific integrated circuit (ASIC) or field-programmable gate array (FPGA)). In general, a heterogeneous computing platform consists of processors with different instruction set architectures (ISAs)." Source: [http://en.wikipedia.org/wiki/Heterogeneous\\_computing](http://en.wikipedia.org/wiki/Heterogeneous_computing).

logical core for processing, see Boyd (2008). Applications include graphics, video games, and image processing.<sup>2</sup> The video card technology has experienced a rapid development mainly driven by 3D rendering (with notable applications in the video game industry), see Boyd (2008). Table 1 reports the GigaFLOPS (i.e., billions of floating point operations per second) for the GPU and the CPU used in this study.

TABLE 1. *Performance of the GPU and the CPU used in this study.*

Type of video card	Stream Cores	Single-precision GFLOPS
NVIDIA GTS 450	192	601.34
Intel Core i7-920	4	31.65

In February 2012, Microsoft released C++ AMP as an extension of Visual Studio 2012 which allows to use the massively parallel and heterogeneous hardware available nowadays. Microsoft is not the only entity active in parallel computation, there are at least two other approaches: the Compute Unified Device Architecture (CUDA) of NVIDIA and the Open Computing Language (OpenCL) of the Khronos Group.

The CUDA (or, more precisely, C for CUDA) approach is an extension of C and although very fast, as documented in Aldrich et al. (2011), it suffers from vendor lock-in, since it only works with NVIDIA GPUs and cannot be used with, for instance, ATI GPUs.<sup>3</sup> OpenCL is designed for heterogeneous parallel computing and since it is not tied to any specific video card manufacturer it can work on NVIDIA and ATI GPUs. The main disadvantage thereof is that OpenCL is based on C rather than C++ and, consequently, involves significantly more manual intervention and low-level programming, see Aldrich et al. (2011) for a discussion.

C++ AMP, in contrast, simply requires a generic GPU and it abstracts the accelerators for the user. In other words, when using C++ AMP, the user does not need to worry about the specific kind of accelerator available in the system. More importantly, the programmer does not need to know or use manual memory management (different for each GPU), since it is performed automatically. This allows to write more natural and high-level programs.

Another big advantage of C++ AMP is its flexibility, boosting developer productivity. It is not necessary to write two different implementations of the same program: one for the NVIDIA GPUs (e.g., using CUDA) and another one for the AMD GPUs (e.g., using OpenCL). It is enough to write one general implementation, since C++ AMP automatically adapts to the

<sup>2</sup>See <http://www.gregcons.com/CppAmp/OverviewAndCppAMPApproach.pdf>.

<sup>3</sup>While Thrust, "a parallel algorithms library which resembles the C++ Standard Template Library (STL)," improves upon the low-level approach of C for CUDA, it still suffers from the same vendor lock-in.

specific hardware available on the target machine. In the case of multiple GPUs in the system (e.g., one integrated within the CPU system and another one discrete), they can be used simultaneously, even if they come from different vendors (e.g., an ATI GPU, an Intel GPU, and an NVIDIA GPU).

To date, the adoption of GPU computing technology in economics and econometrics has been relatively slow compared to other fields. For a literature review, see Morozov and Mathur (2011). This fact can be explained by the need to learn a new syntax, CUDA or OpenCL, the difficulty involved in implementing a different program for each different device, and the more error-prone programming involved in lower-level solutions. We think that C++ AMP is a solution to these problems and that it can help to promote and spread parallel programming in particular and heterogeneous programming in general in the economics and econometrics community.

To show the potential of this approach we replicate the exercise of Aldrich et al. (2011), which uses the value function iteration (VFI henceforth), an algorithm easy to express as a data parallel computational problem, to solve a dynamic equilibrium model.

We find that using VFI with binary search the (optimized) CUDA program is slower than the naive (unoptimized) implementation using C++ AMP. C++ AMP is also faster in a grid search with a Howard improvement step. In both cases the C++ AMP program is almost 3 times faster than the CUDA one.

The rest of the paper is organized as follows. Section 2 provides a simple example that shows the simplicity of the proposed approach compared to CUDA. Section 3 describes the basic idea of parallelization and heterogeneous programming applied to VFI. Section 4 presents the RBC model used in Aldrich et al. (2011). Section 5 reports the time comparison between the two approaches. Section 6 concludes. Appendix A describes a possible optimization approach for C++ AMP.

## 2. A SIMPLE EXAMPLE

This example is intended to show how a low-level programming approach, such as CUDA, can pose a barrier to the widespread adoption of GPU technology in the economics and econometrics community – and how a high-level C++ AMP approach solves the problem.

Let us consider a simple case of matrix multiplication  $\mathbf{C} = \mathbf{AB}$ , where the input matrices  $\mathbf{A}$  and  $\mathbf{B}$  are conformable, and  $\mathbf{C}$  is the output matrix. Let us only focus on the memory management part and compare the source codes given in Listing 1 (CUDA) and Listing 2 (C++ AMP). In each listing, the views `d_A`, `d_B`, and `d_C` (available on the GPU) alias the data in  $\mathbf{A}$ ,  $\mathbf{B}$ , and  $\mathbf{C}$  (available on the CPU), respectively. While the intent and the meaning of the C++ AMP code are relatively clear, the CUDA code involves a significant amount of low-level detail, such as manual memory

allocation and deallocation, manual memory access through pointers (representing values stored in the computer memory using their addresses), manual memory transfer (which also requires manually specifying the memory addresses using pointers and the direction of the transfer), and even such low-level details as the size of a single-precision floating-point data type `float` using the `sizeof` operator. This level of complexity is not only completely unnecessary in the age of modern compilers well capable of automatic type inference, but also more error-prone.<sup>4</sup>

```

1 // "A" and "B" are square matrices
2 // "size" is the row dimension of each matrix
3
4 // allocation
5 cudaError_t err;
6 float *d_A, *d_B, *d_C;
7 err = cudaMalloc(&d_A, size * size * sizeof(float)); //
  input matrix
8 err = cudaMemcpy(d_A, A, size * size * sizeof(float),
  cudaMemcpyHostToDevice);
9 err = cudaMalloc(&d_B, size * size * sizeof(float)); //
  input matrix
10 err = cudaMemcpy(d_B, B, size * size * sizeof(float),
  cudaMemcpyHostToDevice);
11 err = cudaMalloc(&d_C, size * size * sizeof(float)); //
  output matrix
12
13 // deallocation
14 err = cudaMemcpy(C, d_C, size * size * sizeof(float),
  cudaMemcpyDeviceToHost);
15 err = cudaFree(d_A);
16 err = cudaFree(d_B);
17 err = cudaFree(d_C);

```

LISTING 1. CUDA sample

```

1 // "A" and "B" are square matrices
2 // "size" is the row dimension of each matrix
3 array_view<const float, 2> d_A(size, size, A); // input
  matrix
4 array_view<const float, 2> d_B(size, size, B); // input
  matrix
5 array_view<float, 2> d_C(size, size, C); // output matrix
6 d_C.discard_data();

```

LISTING 2. C++ AMP sample

<sup>4</sup>For instance, a programmer using C for CUDA has to remember to deallocate the previously allocated memory – failure to do so may result in memory leaks; in contrast, C++ AMP follows the common C++ approach of relying on the Resource Acquisition Is Initialization (RAII) technique for automatic resource management, see Stroustrup (2000, Section 14.4).

For further details, see “C++ AMP for the CUDA Programmer” by Steve Deitz,<sup>5</sup> which is also the source of the preceding example.

### 3. PARALLELIZATION AND HETEROGENEOUS COMPUTING

In this paper we closely follow the algorithm of Aldrich et al. (2011), providing a simple parallelization scheme for GPUs (here, generalized for an arbitrary accelerator) for the VFI:

- (1) Get the number of processors available,  $P$ , in the accelerator.
- (2) Set a number of grid points,  $N$ , apportioning the state space such that  $N = N_k \times N_z$ , assigning  $N_k$  points to capital and  $N_z$  points to productivity.
- (3) Allot  $N$  grid points to each of the  $P$  processors of the accelerator.
- (4) Set  $V^0$  to an initial guess.
- (5) Copy  $V^0$  to the memory of the accelerator.
- (6) Compute  $V^{i+1}$ , given  $V^i$ , using each processor for its allotted share.<sup>6</sup>
- (7) Repeat step 6 until convergence:  $\|V^{i+1} - V^i\| < \varepsilon$
- (8) Copy  $V^i$  from the accelerator memory to the main memory.

The practical coding is easier compared to CUDA or OpenCL.<sup>7</sup> For example, using CUDA and OpenCL one needs to spend quite a bit of time learning the details of manual memory management of the GPU. Since those are specific to each architecture, this can be a non-trivial problem, see Morozov and Mathur (2011). In C++ AMP memory management is automatic, and so the user has no need to learn the details of each specific architecture.

Step one of the algorithm, which requires the number of processors in the GPUs, is automatically handled in C++ AMP. The same applies for step two, the division of the  $N$  grid points among the  $P$  processors of the GPUs, which is a tuning parameter in CUDA (see Morozov and Mathur (2011) for a nice and thoughtful discussion), which is automatically handled with our approach.

### 4. AN APPLICATION: RBC MODEL REVISITED

As our illustrative example we use the basic RBC model studied in Aldrich et al. (2011). In this model a representative household maximizes the utility function choosing consumption  $\{c_t\}_{t=0}^{\infty}$  and capital  $\{k_t\}_{t=0}^{\infty}$ :

$$\max_{\{c_t\}_{t=0}^{\infty}} \mathbf{E}_0 \left[ \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\eta}}{1-\eta} \right], \quad (1)$$

<sup>5</sup><http://blogs.msdn.com/b/nativeconcurrency/archive/2012/04/11/c-amp-for-the-cuda-programmer.aspx>

<sup>6</sup>Shared memory access ensures that upon termination of this step each processor can read  $V^{i+1}$ .

<sup>7</sup>The source code for our application is available from the authors upon request and will be made available under an open-source license.

where  $\mathbf{E}_0$  denotes the conditional expectation operator,  $\beta$  the discount factor, and  $\eta$  risk aversion. The budget constraint is:

$$c_t + i_t = \omega_t + r_t k_t, \quad (2)$$

with  $\omega_t$  being the wage paid for the unit of labor that the household (inelastically) supplies to the market,  $r_t$  the rental rate of capital, and  $i_t$  the investment. The law of motion for capital accumulation is:

$$k_{t+1} = (1 - \delta)k_t + i_t, \quad (3)$$

where  $\delta$  is the depreciation factor. The technology of a representative firm is  $y_t = z_t k_t^\alpha$ , where productivity  $z_t$  follows an AR(1) process in logs:

$$\log z_t = \rho \log z_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim \mathbf{N}(0, \sigma^2). \quad (4)$$

The resource constraint of the economy is

$$k_{t+1} + c_t = z_t k_t^\alpha + (1 - \delta)k_t. \quad (5)$$

Given that the welfare theorems hold in this economy, we focus on solving the social planner's problem, which can be equivalently stated in terms of a value function  $V$  and its Bellman equation

$$V(k, z) = \max_c \left\{ \frac{c^{1-\eta}}{1-\eta} + \beta \mathbf{E}[V(k', z')|z] \right\} \quad (6)$$

*s.t.*  $k' = z k^\alpha + (1 - \delta)k - c,$

which can be solved with VFI, which is straightforward to parallelize.

For benchmarking purposes we use the same parameter values as in Aldrich et al. (2011):  $\alpha = 0.35$ ,  $\beta = 0.984$ ,  $\delta = 0.01$ ,  $\eta = 2$ ,  $\rho = 0.95$ , and  $\sigma = 0.005$ .

## 5. RESULTS

To implement the C++ AMP program that solves equation (6) we use Microsoft Visual Studio 2012 RC. The CUDA code that solves the same problem, kindly provided at the following location, <http://www.ealdrich.com/Research/GPUVFI/>, is compiled using Visual Studio 2010 since the CUDA code cannot yet be compiled on Visual Studio 2012 RC.

Our test machine is a generic PC with Intel Core i7-920 CPU and NVIDIA GeForce GTS 450 which is an entry-level video card, with a total of 192 stream cores.<sup>8</sup> The test machine runs Windows Server 2012 RC, for specifications see also Table 1.

We run two experiments. The first one uses VFI with binary search, see Heer and Maussner (2005) for a discussion of this algorithm. In this experiment the number of capital points,  $N_k$ , is increased proportionally until a grid of 262,144 points is reached. The productivity process is discretized using the procedure of Tauchen (1986). We report the memory allocation

<sup>8</sup>In a previous version of the paper we had another test machine a standard laptop with Intel Core i7-2630QM and a GeForce 520M with 48 stream cores.

and solution time of the CUDA and the C++ AMP programs in seconds, see Table 2. The Table also reports the total execution time for the CPU.

In the second experiment we use the Howard improvement method and grid search; as pointed out by Aldrich et al. (2011), this algorithm yields a lower return to parallelization when using CUDA. In this case the number of capital points,  $N_k$ , is also increased proportionally until a grid of 262,144 points is reached. We run their experiment with the value function maximized every  $n$ -th iteration of the algorithm, where  $n$  is a tuning parameter decided by the user, see Tables 3 and 4.

We start with the utility of the representative household in the deterministic steady state as our  $V^0$  and the convergence criterion is  $\|V^{i+1} - V^i\| < (1 - \beta) \times 10^{-8}$ , since we are working in double precision. We will come back to this point in Section 6.

Table 2 reports the results for the binary search algorithm with an increasing number of capital points.

TABLE 2. *Observed times (in seconds) for Binary Search. Also reported are the memory allocation (Mem-Alloc) time and the solution time for the CUDA and the C++ AMP programs. For CPU only the total execution time is reported. GPU GeForce GTS 450, CPU Core i7-920.*

$N_k$		32	64	128	256	512	1,024
CUDA (GPU)	Mem-Alloc	1.62	1.62	1.62	1.62	1.62	1.62
CUDA (GPU)	Solution	0.19	0.28	0.25	0.37	0.62	1.22
C++ AMP (GPU)	Mem-Alloc	0.06	0.06	0.06	0.06	0.06	0.06
C++ AMP (GPU)	Solution	1.93	1.72	1.72	1.89	1.96	2.34
C++ (CPU)	Total	0.10	0.24	0.51	1.16	3.04	8.00
$N_k$		2,048	4,096	8,192	16,384	32,768	65,536
CUDA (GPU)	Mem-Alloc	1.62	1.62	1.62	1.62	1.62	1.62
CUDA (GPU)	Solution	2.80	5.10	10.78	22.82	48.67	103.42
C++ AMP (GPU)	Mem-Alloc	0.06	0.06	0.06	0.06	0.06	0.06
C++ AMP (GPU)	Solution	2.72	3.67	5.86	10.44	20.39	41.81
C++ (CPU)	Total	23.31	80.32	288.71	1,477.25	5,692.20	13,930.8

The main result of Table 2 is that the optimized CUDA program (exploiting memory access pattern, called tiling or blocking, to improve performance) is slower than the naive (unoptimized) C++ AMP program on the GeForce GTS 450 GPU.<sup>9</sup> Two more results have to be emphasized. First, the memory allocation step is much faster in the C++ AMP program than in the CUDA one. Second, for small grid the CUDA program is faster than the C++ AMP one. This advantage is lost as soon as the grid increases and in the end the C++ AMP program is almost 3 times faster than the CUDA one. Hence, we conclude that C++ AMP might offer better scalability for larger computational problems.

<sup>9</sup>In a previous version of the paper we also find that the single-precision CUDA program, in this particular case, is slightly faster than the single-precision C++ AMP program.

It is worth noting that C++ AMP can be used in two different ways. The first one is the naive implementation which is very easy and quite similar to traditional C++ code and is the programming approach used in this paper. The second one is the tiling (blocking) optimization, as in the CUDA version, which is more complex, but has the potential to be significantly faster than the naive implementation for specific problems, see Appendix A.

In Tables 3 and 4 we report the results for the grid search algorithm with a Howard step. The value function is maximized only every  $n$ -th iteration of the algorithm. In our experiment we set  $n = \{10, 20\}$ , see Tables 3 and 4, respectively.

TABLE 3. *Observed times (in seconds) for grid search algorithm, with a Howard step every 10 iterations. Also reported are the memory allocation (Mem-Alloc) time and the solution time for the CUDA and the C++ AMP programs. For CPU only the total execution time is reported. GPU GeForce GTS 450, CPU Core i7-920.*

$N_k$		32	64	128	256	512	1,024
CUDA (GPU)	Mem-Alloc	1.62	1.62	1.62	1.62	1.62	1.62
CUDA (GPU)	Solution	0.16	0.17	1.90	0.37	1.11	3.80
C++ AMP (GPU)	Mem-Alloc	0.06	0.06	0.06	0.06	0.06	0.06
C++ AMP (GPU)	Solution	1.81	1.92	1.67	1.79	2.42	4.11
C++ (CPU)	Total	0.04	0.09	0.22	0.61	1.70	7.48
$N_k$		2,048	4,096	8,192	16,384	32,768	65,536
CUDA (GPU)	Mem-Alloc	1.62	1.62	1.62	1.62	1.62	1.62
CUDA (GPU)	Solution	14.32	55.79	220.38	879.56	3,503.06	13,593.30
C++ AMP (GPU)	Mem-Alloc	0.06	0.06	0.06	0.06	0.06	0.06
C++ AMP (GPU)	Solution	8.93	28.52	109.17	365.93	1,438.98	5,702.69
C++ (CPU)	Total	24.87	91.42	440.68	1,743.93	7,137.80	27,239.2

As in the binary search case the naive implementation of C++ AMP is very fast, almost 3 times faster than the CUDA implementation.<sup>10</sup> As explained in Aldrich et al. (2011) the CPU approach is much slower than the GPU one.

<sup>10</sup>In a previous version of the paper, working in single precision, we also find that the C++ AMP program in this case is more than 5 times faster than the CUDA one.

TABLE 4. *Observed times (in seconds) for grid search algorithm, with a Howard step every 20 iterations. Also reported are the memory allocation (Mem-Alloc) time and the solution time for the CUDA and the C++ AMP programs. For CPU only the total execution time is reported. GPU GeForce GTS 450, CPU Core i7-920.*

$N_k$		32	64	128	256	512	1,024
CUDA (GPU)	Mem-Alloc	1.62	1.62	1.62	1.62	1.62	1.62
CUDA (GPU)	Solution	0.12	0.13	0.15	0.26	0.71	2.07
C++ AMP (GPU)	Mem-Alloc	0.06	0.06	0.06	0.06	0.06	0.06
C++ AMP (GPU)	Solution	1.88	1.83	1.69	1.752	2.21	3.10
C++ (CPU)	Total	0.03	0.07	0.16	0.45	1.25	3.91
$N_k$		2,048	4,096	8,192	16,384	32,768	65,536
CUDA (GPU)	Mem-Alloc	1.62	1.62	1.62	1.62	1.62	1.62
CUDA (GPU)	Solution	7.57	28.75	112.89	446.97	1,787.62	7,108.05
C++ AMP (GPU)	Mem-Alloc	0.06	0.06	0.06	0.06	0.06	0.06
C++ AMP (GPU)	Solution	5.62	14.64	49.59	187.93	735.45	2,907.54
C++ (CPU)	Total	14.67	52.74	216.98	930.62	3,934.37	14,913.3

## 6. CONCLUSION

We propose and present a new approach for massively parallel programming in economics, C++ AMP. We show that this approach has a lot of potential. First, the programming simplicity, second, the hardware generality (C++ AMP adapts to different GPUs automatically), and third, the performance (in the considered cases it is much faster than a specialized approach such as CUDA).

As a future research project we are also interested in investigating the reasons behind the speed-up observed for the C++ AMP program in select cases. A potential solution would involve comparing the generated PTX (Parallel Thread Execution) code on the CUDA side with the HLSL (High Level Shader Language) code on the C++ AMP side. However, the current (release candidate) version of C++ AMP does not offer the capability to view the generated HLSL code.

Our results are solely intended as an illustration of a lower bound of this technology for two reasons. First, it is still in the early testing phase and a lot of improvements are expected with the final release by the end of 2012. Second, since C++ AMP is an open standard, it can be implemented by different vendors, such as ATI, for different platforms, such as Unix-like systems.

At the moment there is still a drawback to using double precision computation. For the GPUs used in this study, C++ AMP offers full support for double precision computation on Windows NT 6.2 (i.e., “Windows 8” and “Windows Server 2012”) only. The upcoming release of new drivers from the major vendors, such as NVIDIA and ATI, should solve this problem for the current Windows operating systems (NT 6.1 – i.e., “Windows

7” and “Windows Server 2008 R2”).<sup>11</sup> For the other high-end GPUs, such as GeForce GTX 460, full double precision support is already available. To test the reliability of the C++ AMP program’s results in all of the experiments, we carried out a comparison between the CUDA, the C++ AMP, and the CPU output, finding no differences.

Since heterogeneous computing is a fast-evolving field we expect that additional findings will be forthcoming soon.

---

<sup>11</sup>For more details, see <http://blogs.msdn.com/b/nativeconcurrency/archive/2012/02/07/double-precision-support-in-c-amp.aspx>.

## APPENDIX A. TILING IN C++ AMP

This Appendix briefly describes tiling in C++ AMP.<sup>12</sup> Tiling is a way to organize computation in order to use the specific memory areas more efficiently, see McKellar and Coffman (1969).

In general, modern GPUs have (at least) two types of memory areas – the *global memory* and the *local memory* (also called "shared memory" – and called "tile\_static memory" in C++ AMP). Accessing the *global memory* is much slower than the *local memory*.

Unfortunately managing the *local memory* efficiently is neither trivial nor automatic. The C++ AMP tiled model allows to manage the access to and to organize the use of the *local memory* manually. The C++ AMP compiler does a mechanical transformation behind the scenes, always resulting in a tiled version of the program, even when the user is not using tiling explicitly (such as when using the implicitly tiled simple model, as we do in this study). When the user needs more performance it is possible to go beyond what the compiler can do automatically by writing the explicitly tiled version of the program, deciding how to organize (partition) the use of and the access to the different kinds of memory areas manually.

In fact, in some cases tiled computations have been shown to be twice as fast as their non-tiled counterparts. These performance improvements are strongly dependent on the problem at hand. It is usually advisable to use tiling in the case of a large number of repeated operations, such as matrix inversion for large matrices.

The current Visual Studio 2012 is still in the Release Candidate stage and it is possible that the further optimizations performed by the compiler will close the speed gap between the tiled and the non-tiled implementations in the future.

## REFERENCES

- Aldrich, E. M., Fernández-Villaverde, J., Gallant, A. R., and Rubio Ramirez, J. F. (2011). Tapping the Supercomputer Under Your Desk: Solving Dynamic Equilibrium Models with Graphics Processors. *Journal of Economic Dynamics and Control*, 35:386–393.
- Boyd, C. (2008). Data-parallel computing. *Queue*, 6(2):30–39.
- Heer, B. and Maussner, A. (2005). *Dynamic General Equilibrium Modelling: Computational Methods and Applications*. Springer, Berlin.
- McKellar, A. C. and Coffman, E. G. (1969). Organizing matrices and matrix operations for paged memory systems. *Communications of the ACM*, 12:153–165.
- Morozov, S. and Mathur, S. (2011). Massively Parallel Computation Using Graphics Processors with Application to Optimal Experimentation in Dynamic Control. *Computational Economics*, pages 1–32.

<sup>12</sup>For more details, see <http://msdn.microsoft.com/en-us/magazine/hh882447.aspx>.

Stroustrup, B. (2000). *The C++ Programming Language*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 3rd edition.

Tauchen, G. (1986). Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions. *Economics Letters*, 20:177 – 181.

- 2011-52: Lars Stentoft: What we can learn from pricing 139,879 Individual Stock Options
- 2011-53: Kim Christensen, Mark Podolskij and Mathias Vetter: On covariation estimation for multivariate continuous Itô semimartingales with noise in non-synchronous observation schemes
- 2012-01: Matei Demetrescu and Robinson Kruse: The Power of Unit Root Tests Against Nonlinear Local Alternatives
- 2012-02: Matias D. Cattaneo, Michael Jansson and Whitney K. Newey: Alternative Asymptotics and the Partially Linear Model with Many Regressors
- 2012-03: Matt P. Dziubinski: Conditionally-Uniform Feasible Grid Search Algorithm
- 2012-04: Jeroen V.K. Rombouts, Lars Stentoft and Francesco Violante: The Value of Multivariate Model Sophistication: An Application to pricing Dow Jones Industrial Average options
- 2012-05: Anders Bredahl Kock: On the Oracle Property of the Adaptive LASSO in Stationary and Nonstationary Autoregressions
- 2012-06: Christian Bach and Matt P. Dziubinski: Commodity derivatives pricing with inventory effects
- 2012-07: Cristina Amado and Timo Teräsvirta: Modelling Changes in the Unconditional Variance of Long Stock Return Series
- 2012-08: Anne Opschoor, Michel van der Wel, Dick van Dijk and Nick Taylor: On the Effects of Private Information on Volatility
- 2012-09: Annastiina Silvennoinen and Timo Teräsvirta: Modelling conditional correlations of asset returns: A smooth transition approach
- 2012-10: Peter Exterkate: Model Selection in Kernel Ridge Regression
- 2012-11: Torben G. Andersen, Nicola Fusari and Viktor Todorov: Parametric Inference and Dynamic State Recovery from Option Panels
- 2012-12: Mark Podolskij and Katrin Wasmuth: Goodness-of-fit testing for fractional diffusions
- 2012-13: Almut E. D. Veraart and Luitgard A. M. Veraart: Modelling electricity day-ahead prices by multivariate Lévy
- 2012-14: Niels Haldrup, Robinson Kruse, Timo Teräsvirta and Rasmus T. Varneskov: Unit roots, nonlinearities and structural breaks
- 2012-15: Matt P. Dziubinski and Stefano Grassi: Heterogeneous Computing in Economics: A Simplified Approach