

Cheela, Bhagath; DeHon, André; Fernández-Villaverde, Jesús; Peri, Alessandro

Article

Programming FPGAs for economics: An introduction to electrical engineering economics

Quantitative Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Cheela, Bhagath; DeHon, André; Fernández-Villaverde, Jesús; Peri, Alessandro (2025) : Programming FPGAs for economics: An introduction to electrical engineering economics, Quantitative Economics, ISSN 1759-7331, The Econometric Society, New Haven, CT, Vol. 16, Iss. 1, pp. 49-87,
<https://doi.org/10.3982/QE2344>

This Version is available at:

<https://hdl.handle.net/10419/320332>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/4.0/>

Programming FPGAs for economics: An introduction to electrical engineering economics

BHAGATH CHEELA

Department of Electrical and Systems Engineering, University of Pennsylvania

ANDRÉ DEHON

Department of Electrical and Systems Engineering, University of Pennsylvania

JESÚS FERNÁNDEZ-VILLAYERDE

Department of Economics, University of Pennsylvania

ALESSANDRO PERI

Department of Economics, University of Colorado, Boulder

We show how to use field-programmable gate arrays (FPGAs) and their associated high-level synthesis (HLS) compilers to solve heterogeneous agent models with incomplete markets and aggregate uncertainty (Krusell and Smith (1998)). We document that the acceleration delivered by one single FPGA is comparable to that provided by using 69 CPU cores in a conventional cluster. The time to solve 1200 versions of the model drops from 8 hours to 7 minutes, illustrating a great potential for structural estimation. We describe how to achieve multiple acceleration opportunities—pipeline, data-level parallelism, and data precision—with minimal modification of the C/C++ code written for a traditional sequential processor, which we then deploy on FPGAs easily available at Amazon Web Services. We quantify the speedup and cost of these accelerations. Our paper is the first step toward a new field, electrical engineering economics, focused on designing computational accelerators for economics to tackle challenging quantitative models. Replication code is available on [Github](#).

Bhagath Cheela: cheelabhagath@gmail.com

André DeHon: andre@acm.org

Jesús Fernández-Villaverde: jesusfv@econ.upenn.edu

Alessandro Peri: alessandro.peri@colorado.edu

We thank Yicheng Li (UPenn, Engineering) for the initial implementation of our hardware design. We thank Lucas Ladenburger, Marina Leah McCann, and Paro Suh (CU Boulder, Economics) for outstanding research assistance. We thank Andrew Monaghan and the CU Boulder Research Computing Center for providing valuable insights. We also thank Giuseppe Bruno and Riccardo Russo (Bank of Italy) for testing an early version of the tutorial associated with this paper and Victor Duarte, Mahdi E. Kahou, and Jesse Perla for their comments. This project used: (i) the RMACC Summit supercomputer, supported by the National Science Foundation (awards ACI-1532235 and ACI-1532236), CU Boulder and Colorado State University; (ii) AWS Credits awarded under the NSF CC* Hybrid Cloud Award OAC-1925766, Research Computing, CU Boulder, 2022. This project was also supported by the Undergraduate Research Experiences for Diversity Grant, 2021, Institute of Behavioral Science, University of Colorado, USA.

KEYWORDS. FPGA acceleration, heterogeneous agents, aggregate uncertainty, electrical engineering economics, cloud computing.

JEL CLASSIFICATION. C6, C63, C88, D52.

1. INTRODUCTION

Computations play a crucial role in nearly all fields of economics. The over 3000 pages of the four volumes of the renowned *Handbook of Computational Economics* (Amman et al. (2018)) demonstrate this point. Interestingly, economists have paid much less attention to hardware. Except for a few papers (among others, Nagurney (1996), Aldrich, Fernández-Villaverde, Gallant, and Rubio-Ramírez (2011), Fernández-Villaverde and Valencia (2018), Duarte, Duarte, Fonseca, and Montecinos (2019), and Peri (2020)), researchers have taken hardware as a given.

This lack of interest is unfortunate. While the speeds of single-core central processing units (CPU) are stalling, specialized accelerators (i.e., hardware designed to perform specific operations) continue to deliver greater performance at inexpensive price points. Since advanced computation is playing a growing role in all industries and research, the development and programming of customized accelerator chips have become cheaper and more accessible.

Fields like biology, genetics, medicine, and physics have been at the vanguard of this process. Over the years, research in these fields has seen widespread adoption of a chip whose hardware can be programmed to solve application-specific algorithms: field-programmable gate arrays (FPGA). FPGAs have been successfully deployed to accelerate a variety of applications: DNA matching (Hoang (1993)), molecular dynamics (Azizi, Kuon, Egier, Darabiha, and Chow (2004)), Basic Local Alignment Search Tool (Herbordt, Model, Gu, Sukhwani, and VanCourt (2006)), astrophysics particle simulator (Berczik et al. (2009)), and cancer treatment (Young-Schultz, Lilge, Brown, and Betz (2020)), among others.

So far, the skills required to implement algorithms in hardware have represented a significant entry barrier, particularly in disciplines characterized by small research teams, such as economics. These domain-specific knowledge requirements are apparent in Peri (2020), the only other attempt we are aware of employing FPGA technology in economics. Following the industry standard for chip design, Peri (2020) employs a hardware description language to describe a circuit at the register transfer level (RTL) that solves a Bellman equation via value function iteration on an FPGA. Much like working with assembly language, this RTL design approach uses an intricate low-level syntax to configure the hardware, rendering the code hardly accessible to economists. More importantly, it requires a working knowledge of digital design, a subject that most economists are unfamiliar with. Our paper describes how to reduce these barriers for economics.

How do we get higher performance from a chip while retaining a reasonable easiness of programmability? By replacing the lower-level RTL design in Peri (2020) (the RTL approach) with the more accessible FPGA HLS C-to-gates compilers (the HLS ap-

proach).¹ Here, we follow a long history of increasingly sophisticated automation in computer science to map higher-level abstractions down to low-level implementations. Our HLS approach allows economists to implement complex circuits with minimal to no knowledge of hardware design and delegate the technical details to the compiler. More significantly, our approach dramatically reduces the coding and debugging time compared to RTL, facilitating the implementation of far more complicated dynamic models and the exploration of different model features, both key aspects of practical research. We illustrate these ideas by accelerating a major workhorse model in economics: the incomplete markets, heterogeneous agent model with aggregate uncertainty of [Krusell and Smith \(1998\)](#).

We introduce the concepts of FPGAs' programming and optimizations in the context of a simple yet illustrative example: the accumulation of array elements. Here, we intuitively showcase *all* the hardware design techniques deployed to accelerate the [Krusell and Smith \(1998\)](#) algorithm, illustrating how progressive code modifications translate into gradually higher-performance FPGA circuits. Along the way, we compile a list of hardware issues that economists are likely to encounter when accelerating their applications, including all the performance bottlenecks faced while accelerating the [Krusell and Smith \(1998\)](#) algorithm. This approach allows us to distill essential hardware design principles for accelerating dynamic models without necessitating economists to become hardware designers themselves. In the process, we highlight the remarkable development advantages over the RTL approach proposed in [Peri \(2020\)](#) in terms of accessibility, coding, and debugging time. Last, we discuss the portability of our code, including, among other aspects, how `Matlab 2022a` and later editions can directly implement the accumulator problem.

We deploy our HLS approach on Amazon Web Services (AWS) to accelerate the solution of an incomplete markets, heterogeneous agent model with aggregate uncertainty. After the pioneering work of [Krusell and Smith \(1998\)](#), heterogeneous agent models have been used to study business cycle fluctuations, monetary and fiscal policy, climate change, life-cycle decisions, industry dynamics, and international trade. Also, there has been tremendous interest in the development of solution methods well suited for these models, such as [Algan, Allais, and Den Haan \(2008\)](#), [Reiter \(2009\)](#), [Den Haan and Rendahl \(2010\)](#), [Maliar, Maliar, and Valli \(2010\)](#), [Reiter \(2010\)](#), [Young \(2010\)](#), [Algan, Allais, Den Haan, and Rendahl \(2014\)](#), [Pröhl \(2015\)](#), [Achdou, Han, Lasry, Lions, and Moll \(2021\)](#), [Bhandari, Evans, Golosov, and Sargent \(2017\)](#), [Brumm and Scheidegger \(2017\)](#), [Judd, Maliar, Maliar, and Tsener \(2017\)](#), [Bayer and Luetticke \(2018\)](#), [Childers \(2018\)](#), [Mertens and Judd \(2018\)](#), [Winberry \(2018\)](#), [Fernández-Villaverde, Hurtado, and Nuño \(2019\)](#), [Auclet, Rognlie, and Straub \(2020\)](#), [Bilal \(2021\)](#), and [Kahou, Fernández-Villaverde, Perla, and Sood \(2021\)](#), among many others.

More concretely, we work with the canonical heterogeneous agent model in [Den Haan, Judd, and Juillard \(2010\)](#). The authors proposed this economy as a computational suite to test the accuracy of solution methods for heterogeneous agent models precisely

¹From the earliest demonstrations ([Babb et al. \(1999\)](#)) to the launch of the first commercial compilers (`Streams-C`, [Frigo, Gokhale, and Lavenier \(2001\)](#), [Snider \(2002\)](#)), FPGA compilers have steadily improved, making the programming of FPGAs cost-effective in terms of coding time.

because of its canonicity. We solve the model employing the [Krusell and Smith \(1998\)](#) algorithm as implemented by [Maliar, Maliar, and Valli \(2010\)](#). We pick this solution algorithm because of its simplicity and accuracy. A further advantage of this code is that it was not written to extract performance from custom accelerators, limiting the possible biases in our analysis. We code our solution in C/C++, the fastest programming environment for computation in economics ([Aruoba and Fernández-Villaverde \(2015\)](#)).² For the FPGA, we employ the industry gold standard Xilinx high-level synthesis (HLS) compilers using Vitis HLS (v2021.2, 64-bit) and the OpenCL interface ([Xilinx, Inc. \(2020\)](#)). For the CPU, we use the G++ 9.4.0 and mpic++ 4.1.1 (Open MPI) compilers, which implement state-of-the-art sequential and parallel execution run-time optimizations (and whose importance will be clear momentarily).³

Our first exercise runs the code in one FPGA and one CPU core. The code is *exactly* the same in both cases except when we apply the `#pragma` directives available to the FPGA designer. The `#pragmas` instruct the HLS compiler on how to *design* the FPGA hardware, that is, on how to connect *physical* logical resources in the FPGA to exploit the maximum parallelism, pipelining, and data distribution required to solve our algorithm efficiently.⁴ The CPU cannot offer this feature because its hardware is not programmable. The CPU code is compiled using the `-O3` flag, the most aggressive standard-compliant optimization and one that delivers an executable file that is hard to beat via hand-made optimizations, even for highly experienced programmers. Thus, our code compares the best performance offered by one FPGA with the best performance offered by one CPU core, making the comparison meaningful. The FPGA delivers a speedup of nearly 69 times against an Intel Xeon Scalable Processor (Cascade Lake, second generation) core. That is, we solve the same heterogeneous agent model 69 times faster in one FPGA than in one CPU core, reducing the time to solve 1200 economies from 8 hours to 7 minutes.

Our second exercise scales up from one FPGA to eight and from one CPU core to 48 to compare their performance when a researcher has access to multicore acceleration. In this case, we use Open MPI (the de facto standard for parallelization in large clusters of CPUs) to parallelize our code as deployed in the CPUs. In particular, we ask each CPU core to solve the same model many times (we call each solution an “economy”). For example, if we have 48 cores, we ask each CPU core to solve 25 economies for a total of 1200 economies, and we compare the time required against the time that the FPGAs require to solve 1200 economies. This research design is the *best case scenario* for parallelization in CPU cores, as it minimizes the communication across CPU cores and its associated overheads. Other parallelization schemes, such as solving one economy simultaneously in several cores (perhaps more relevant in practice because the model to solve is complex), will deliver worse performance for multicore CPU clusters because

²We code our FPGA kernel functions in C to comply with the FPGA HLS C-to-gates compiler requirement. Similarly, we code our CPU kernel functions in C to avoid the abstraction penalty that may arise from C++ object-oriented programming. Thus, we give the CPU the best fighting chance against FPGA acceleration.

³All replication codes are available on a Github repository: <https://github.com/AleP83/FPGA-Econ.git>.

⁴Our online tutorial Cheela, DeHon, Fernández-Villaverde, and Peri (2023) supplies a detailed guide on how to utilize `#pragmas` to accelerate the solution of our model.

of time lost in data transfers. We find that one FPGA solves 1200 economies 1.48 times faster than 48 CPU cores do and that eight FPGAs solve 1200 economies 549 times faster than one CPU core and 12 times faster than 48 CPU cores.

The FPGA speedups in these exercises are accompanied by large cost savings (less than 20% of the CPU cost) and even more impressive energy savings (less than 6% of the CPU energy consumption), a growing concern due to environmental goals set up by universities and research institutions. An additional attractive feature of FPGAs is that they are easily available either for purchase at economical prices or at commercial cloud services, such as AWS (the service we use in this paper), at low costs per hour.

Next, we inspect the mechanisms that account for the FPGA speedups. First, we use pipelining of complex equations to start a new calculation composed of hundreds of primitive operations in each cycle. In the process, we tune the precision and data representation to achieve efficient pipelining and support additional parallelism. Second, we exploit loop-level data parallelism to perform computations on multiple independent pipelines simultaneously. Third, we employ coarse-grained, data-level parallelism to compute multiple economies in parallel.

While our application is focused on the canonical model in [Den Haan, Judd, and Juillard \(2010\)](#), the acceleration techniques we describe for its solution can be easily generalized for solving more complicated heterogeneous agent models. We have in mind, for example, the classes of HANK ([Kaplan, Moll, and Violante \(2018\)](#), [Bayer, Luetticke, Pham-Dao, and Tjaden \(2019\)](#)) and climate change models ([Cai and Lontzek \(2019\)](#), [Cruz Álvarez and Rossi-Hansberg \(2021\)](#), [Krusell and Smith \(2022\)](#)) that have become so influential in recent years. Both classes of models face a whole new range of computational challenges with respect to the basic framework in [Krusell and Smith \(1998\)](#) that have prevented their full deployment for policy advising. FPGAs are well suited to deal with these larger models. Nonetheless, we prefer to illustrate our approach with the model in [Den Haan, Judd, and Juillard \(2010\)](#) instead of jumping directly to more complex environments for transparency. Over the last 25 years, we have accumulated so much knowledge about what works (and what does not) while solving models à la [Krusell and Smith \(1998\)](#) that the reader can appreciate, more quickly, the novelty of our work, the speed gains, and its accuracy.

Similarly, the extra speed that FPGAs deliver can also be useful when structurally estimating the model using micro and macro data or checking for the robustness of the results with respect to different parameter values. In both cases, we need to solve the model for many parameter values. Thus, speed is a first-order consideration.

In summary, the design of specialized computational accelerators for specific yet vital computational tasks in economics at attractive price points and reasonable programming complexity holds much promise. Programming FPGAs to solve heterogeneous agent models is but a first step into a rich area of research.

The rest of the paper is organized as follows. Section 2 presents the model and its calibration. Section 3 details the solution algorithm. Section 4 introduces the building blocks of FPGAs' programming and optimizations. Section 5 describes the acceleration schemes. Section 6 reports quantitative results and performs robustness tests. Section 7 isolates the acceleration channels responsible for the speed gains. Section 8 concludes.

A Supplemental Appendix (Cheela, DeHon, Fernández-Villaverde, and Peri (2024)) and tutorial (Cheela et al. (2023)) provide further details.

2. THE MODEL

This section presents the incomplete markets model in Den Haan, Judd, and Juillard (2010). This is a production economy with aggregate uncertainty, and a unit mass of infinitely lived ex ante identical households that experience uninsurable idiosyncratic shocks to their employment status and are subject to a borrowing constraint.

The household's problem Households consume $c_t \in \mathbb{R}_+$, supply labor \bar{l} (paid at a wage w_t), and accumulate capital $k_{t+1} \in \mathbf{K} \subseteq \mathbb{R}_+$ (which receives a rental rate r_t). Households suffer idiosyncratic shocks to their employment status $\epsilon_t \in \{0, 1\}$, which equals one if employed and 0 if unemployed. We will specify the stochastic process for ϵ_t below. Also, τ_t is the labor income tax rate, μw_t is the (tax-free) unemployment benefit, and δ is the depreciation rate.⁵

Thus, households choose sequences of consumption and physical capital to solve:

$$\max_{\{c_t, k_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\gamma} - 1}{1-\gamma} \quad (1)$$

$$\text{s.t. } c_t + k_{t+1} = [(1 - \tau_t)\bar{l}\epsilon_t + \mu(1 - \epsilon_t)]w_t + (1 + r_t - \delta)k_t \quad (2)$$

$$k_{t+1} \geq 0. \quad (3)$$

The firm's problem A representative, perfectly competitive firm uses per capita capital $K_t \in \mathbf{M} \subseteq \mathbb{R}_+$ and the employment rate $L_t \in \mathbb{R}_+$ to produce a per capita final good with a technology $Y_t = A_t K_t^\alpha (\bar{l}L_t)^{1-\alpha}$, where $0 < \alpha < 1$. The aggregate productivity A_t follows a two-state Markov process over the support $\mathbf{A} = \{a_b, a_g\}$, where $a_g = (1 + \Delta_A)$ is the good realization and $a_b = (1 - \Delta_A)$ is the bad realization.

Competition in the input and output markets implies that:-

$$r_t = \alpha A_t \left(\frac{\bar{l}L_t}{K_t} \right)^{1-\alpha}, \quad (4)$$

and

$$w_t = (1 - \alpha) A_t \left(\frac{K_t}{\bar{l}L_t} \right)^\alpha. \quad (5)$$

Government The government uses labor income taxes τ_t to finance unemployment benefits (in terms of wages) μ ,

$$\tau_t \bar{l}L_t = \mu(1 - L_t), \quad (6)$$

where $u_t = 1 - L_t$ is the unemployment rate.

⁵Our notation follows Den Haan, Judd, and Juillard (2010), except that we drop the subindex for individual households. We are more explicit about the choice sets than what is standard to facilitate readability. When unemployment benefits are set to zero, the model coincides with the model in Krusell and Smith (1998).

TABLE 1. Calibrated parameters.

α	0.36	Output capital share
β	0.99	Quarterly discount factor
γ	1	Relative risk aversion coefficient
δ	0.025	Quarterly depreciation rate
μ	0.15	Unemployment benefits in terms of wages
\bar{l}	1/0.9	Time endowment
Δ_A	0.01	Aggregate productivity shock size

Aggregate law of motion The cross-sectional distribution of households over capital holdings and employment status, Γ , follows the law of motion:

$$\Gamma_{t+1} = \mathcal{H}(\Gamma_t, A_t, A_{t+1}). \quad (7)$$

Equilibrium Given an exogenous transition law for $\{A, \epsilon\}$, a recursive competitive equilibrium is the set of prices $\{w, r\}$, policy function $k'(\cdot)$, tax rate τ , and law of motion $\mathcal{H}(\cdot)$ for the cross-sectional distribution Γ such that: (i) given the individual household state $\{k, \epsilon; \Gamma, A\}$, prices $\{w, r\}$ and the laws of motion of $\{A, \epsilon\}$ and Γ , the policy function $k'(\cdot)$ solves the Bellman equation representation of the household's sequential problem in (1)–(3); (ii) given $\{\Gamma, A\}$, input factor prices $\{w, r\}$ equal the marginal products (4)–(5); (iii) given A, τ balances the government budget, (6); (iv) the markets for labor and capital clear; (v) given $\{w, r, \Gamma, k', A\}$ and the transition laws for $\{A, \epsilon\}$, the law of motion $\mathcal{H}(\cdot)$ satisfies (7).

Calibration To ensure the maximum comparability of results, we replicate the calibration of Den Haan, Judd, and Juillard (2010). The time unit is a quarter. We discretize the joint transition of aggregate productivity and idiosyncratic employment status using the transition matrix in Table 2 in Den Haan, Judd, and Juillard (2010). The transition probabilities are designed such that the unemployment rate depends only on aggregate productivity, and they take values $u_b = u(1 - \Delta_A)$ in bad times and $u_g = u(1 + \Delta_A)$ in good times, $u_b > u_g$. Table 1 summarizes the rest of the parameters as described in Den Haan, Judd, and Juillard (2010).

We will refer to the set of parameters $\underline{\theta} = \{\alpha, \beta, \gamma, \delta, \mu, \bar{l}, \Delta_A\}$ calibrated as in Table 1 as the *baseline economy*. Without loss of generality, we will refer to a generic $\underline{\theta}$ as an *economy*.

3. THE SOLUTION ALGORITHM

We solve the model using the stochastic simulation algorithm described in Maliar, Maliar, and Valli (2010). Following Krusell and Smith (1998), we assume that households are boundedly rational and perceive that only a finite set of moments of Γ affect future prices. In the numerical exercise, we restrict our attention to the law of motion that describes the evolution of the first moment of the cross-sectional distribution of the per

capita stock of capital, $m \in \mathbf{M}$ among households:

$$m' = H(m, A, A')$$

(notice the change in notation from $\mathcal{H}(\cdot)$ to $H(\cdot)$).

A. Grids We define the intervals on the household's capital $\mathbf{K} \equiv [k_{\min}, k_{\max}] = [0, 1000]$ and first moment of the cross-sectional distribution $\mathbf{M} \equiv [m_{\min}, m_{\max}] = [30, 50]$. We discretize them with $N_k = 100$ and $N_M = 4$ grid points, respectively. Section 6.5 explores alternative grid sizes.

B. Individual households' problem (IHP) A stationary solution to the optimization problem (1) is the saving policy function $k' : \mathbf{K} \times \{0, 1\}_\epsilon \times \mathbf{M} \times \mathbf{A} \rightarrow \mathbb{R}_+$:

$$k' = [\mu(1 - \epsilon) + (1 - \tau)\bar{l}\epsilon]w + (1 - \delta + r)k - \left\{ \lambda + \beta \mathbb{E} \left[\frac{1 - \delta + r'}{((\mu(1 - \epsilon') + (1 - \tau')\bar{l}\epsilon')w' + (1 - \delta + r')k' - k'(k'))^\gamma} \right] \right\}^{-1/\gamma}, \quad (8)$$

where $k' \equiv k'(k, \epsilon, m, A)$, $\lambda \equiv \lambda(k, \epsilon, m, A)$, $k'(k') \equiv k'(k'(k, \epsilon, m, A), \epsilon', m', A')$ and that satisfies the occasionally binding constraint $k' \geq 0$ and complementary slackness condition $\lambda \cdot k' = 0$, with $\lambda \geq 0$.

We find a solution to the IHP using the following iterative Euler equation algorithm:

(i) *Initial guess.* Guess an initial policy function k'_0 . We set $k'_0 \equiv k'_0(k, \epsilon, m, A) = 0.9 \cdot k$.

(ii) *Iteration step.* For each iteration step $i \geq 0$ and given the guess k'_i :

$$\widehat{k}'_{i+1} = \Phi k'_i \quad (\text{a}) \quad \text{Solve (8),}$$

$$k'_{i+1} = \eta_k \widehat{k}'_{i+1} + (1 - \eta_k)k'_i \quad (\text{b}) \quad \text{Update Guess.}$$

(a) For any state $(k, \epsilon, m, A) \in \mathbf{K} \times \{0, 1\}_\epsilon \times \mathbf{M} \times \mathbf{A}$:

- Set the Lagrange multiplier $\lambda(k, \epsilon, m, A) = 0$.
- Substitute the guess k'_i and compute the right-hand side of equation (8).
- Update the left-hand side. If \widehat{k}'_{i+1} falls outside the capital grid set, replace it with the closest boundary of the individual capital grid, $\{k_{\min}, k_{\max}\}$.

(b) After completing step (a), let $\eta_k = 0.7$ and set the $(i + 1)$ -iteration policy function guess to

$$k'_{i+1} = \eta_k \widehat{k}'_{i+1} + (1 - \eta_k)k'_i. \quad (9)$$

(iii) *Convergence criterion.* Repeat the iteration step (ii) until convergence of the policy function in the sup norm:

$$\rho(k'_{i+1}, k'_i) = \max_{(k, \epsilon, m, A) \in \mathbf{K} \times \{0, 1\}_\epsilon \times \mathbf{M} \times \mathbf{A}} |k'_{i+1} - k'_i| < \varepsilon_k = 1e(-8). \quad (10)$$

C. Aggregate law of motion (ALM) Following [Maliar, Maliar, and Valli \(2010\)](#), we parameterize the law of motion of the per capita stock of capital, $m \in \mathbf{M}$ as

$$\ln m' = f(a, m; b) = b_1(a) + b_2(a) \ln m, \quad a \in \{a_b, a_g\}, \quad (11)$$

where the second equality assumes the conditional expectation of $\ln m'$ to be linear in $\ln m$ and aggregate-state dependent.

D. The fixed-point algorithm We estimate the vector of aggregate state-dependent coefficients b (henceforth, ALM coefficients) using a nested fixed-point iterative algorithm:

(i) *Initializations.* Set initial values for $(b_1(a_g), b_2(a_g)) = (b_1(a_b), b_2(a_b)) = \{0, 1\}$. Set the size of the cross-sectional distribution to $J = 10,000$ households and the length of the stochastic simulation to $T = 1100$. Use a pseudo-random number generator to draw the aggregate shocks $\{a_t\}_{t=1}^{1100}$ and the idiosyncratic shocks to the employment status $\{\epsilon_{t,j}\}_{t=1, j=1}^{T=1100, J=10,000}$. Set the initial cross-sectional distribution of the households' capital holdings.⁶ Use equations (4), (5), and (6) to compute w_t , r_t , and τ_t .

(ii) *Iteration step.* For each iteration step $i \geq 0$:

(a) **IHP.** Estimate the households' capital holdings policy functions $k'(k, \epsilon, m, A)$, by solving the IHP.

(b) **Simulation.** At each $t \in \{1, \dots, 1100\}$, given the initial capital holdings distribution, the idiosyncratic and aggregate stochastic processes, and the policy functions:

(i) *Accumulation step.* Compute m_t , the cross-sectional average of households' capital holdings

$$m_t = \frac{1}{J} \sum_{j=1}^J k_{j,t}. \quad (12)$$

(ii) *Interpolation step.* For each household $j \in \{1, \dots, 10,000\}$, use linear interpolation to determine the next period household capital holdings, given the period t idiosyncratic $\{k_{t,j}, \epsilon_{t,j}\}$ and aggregate $\{m_t, A_t\}$ states.⁷

⁶Following [Maliar, Maliar, and Valli \(2010\)](#): (i) we set the initial capital distribution to the steady-state value of capital in a deterministic model with employment rate $L = 1/\bar{l} = 0.9$; (ii) we iteratively update the initial capital distribution (with the capital distribution associated with $T = 1100$) if the metric measuring the convergence of the ALM coefficients in (13) is higher than $1e(-6)$.

⁷[Maliar, Maliar, and Valli \(2010\)](#) use a spline interpolation scheme. Because moving the splines to an FPGA involves some extra work that would distract from the clarity of exposition, we leave the implementation of this feature for future research. In addition, we precompute aggregate and idiosyncratic shocks to ensure comparability across different software. These are the only two differences with their original code, available at <https://lmaliar.ws.gc.cuny.edu/codes/>.

(c) **ALM.** Estimate $\hat{b}^i(a) = (\hat{b}_1^i(a), \hat{b}_2^i(a))$ with $a \in \{a_b, a_g\}$ by running the OLS regression associated with (11), after discarding the first 100 observations, $t = 1, \dots, 100$.

(d) Let $\eta_b = 0.3$. Set the $(i + 1)$ -iteration ALM coefficients to

$$b_l^{i+1}(a) = \eta_b \hat{b}_l^i(a) + (1 - \eta_b) b_l^i(a), \quad l \in \{1, 2\}, a \in \{a_b, a_g\}.$$

(iii) *Convergence criterion.* Repeat the iteration step (ii)(a)–(ii)(d) until convergence of the ALM coefficients in the Euclidean norm:

$$\sqrt{\sum_{l \in \{1, 2\}, a \in \{a_b, a_g\}} (b_l^{i+1}(a) - b_l^i(a))^2} < \varepsilon_b = 1e(-8). \quad (13)$$

4. BUILDING BLOCKS OF FPGAs' OPTIMIZATIONS

Next, we introduce the building blocks of FPGA optimizations: data precision, loop pipelining, loop unrolling, and memory management.

We do so by focusing on a simple example: the sum of eight numbers, $\sum_{i=0}^8 X_i$ (“eight” is just for concreteness, nothing essential depends on it). This accumulation problem, which is equivalent to the problem of computing the cross-sectional average of individual capital holdings in equation (12) by summing J numbers, allows us to showcase all the acceleration techniques deployed in our main application. In the process, we illustrate how these circuit design tools can be exploited with minimal alterations of existing C code and we highlight the development advantages, namely the RTL approach proposed in Peri (2020), underscoring the potential of HLS programming for economics.

For expositional convenience, we focus our discussion on the *kernel*, that is, the algorithm the hardware designer seeks to accelerate by programming the custom logic of the FPGA (the accumulator in this section; the Krusell and Smith (1998) algorithm, later on). Our tutorial (Cheela et al. (2023)) delves into the details of the nonkernel boilerplate code required to implement our application in hardware. Since the kernel is the only code changing across the exercises in this section, this choice comes without loss of generality.

4.1 Programming FPGAs using HLS

Programming an FPGA is the *design* of a physical circuit by connecting the chip's custom logic (CL). Consequently, we will use the term “CL design” interchangeably to refer to circuit design moving forward. FPGA chips are built on the fundamental idea that we can implement Boolean combinational functions using memories, also referred to as lookup tables. By strategically organizing enhanced versions of these lookup tables, known as configurable logic blocks (CLBs), within a programmable grid, FPGAs enable the physical implementation of circuits that would otherwise require months to fabricate on traditional silicon chips.

This premise underscores a conceptual difference between *hardware* and *software* programming. While software programming yields an executable application designed

to run on existing hardware, FPGA programming creates a circuit. Beyond this conceptual difference, software and hardware programming share a historical quest to reduce implementation costs by raising the level of abstraction. By leveraging the use of *compilers*, software programming has progressively replaced low-level (assembly) language with increasingly sophisticated higher-level languages (C/C++, Fortran, Julia, etc.). The early adoption of compilers was driven by the necessity to take advantage of the rapid increase in the number of transistors in the late twentieth century (Moore's law). The stalling speed of single-core central processing units (CPU) has made the *compiler* approach increasingly more attractive in hardware programming, too, as it makes the design of circuits for customized hardware less time-consuming. After decades of effort, commercially supported compilers (Xilinx, Inc. (2020), Corporation Intel (2021)) are now available to replace the widely used low-level register transfer level (RTL) design. Similar to assembly, RTL coding employs a low-level syntax to design the required hardware configurations, making the code hardly accessible to trained economists. Our goal is to explain how FPGA compilers break this barrier.

To make this point, we illustrate the simplicity of implementing the sum of $J = 8$ numbers in hardware using HLS. This operation requires just compiling existing C code written to be executed on the CPU with the HLS `v++` compiler in place of, say, the `g++` compiler,

```
1 void hw_loop(array_type st_k[J], array_type &reduced_sum){
2     array_type sum = 0.;
3     for(int i=0; i<J; i++) {
4         sum+=st_k[i];
5     }
6     reduced_sum = sum;
7     return;
8 }
```

LISTING 1. Sequential accumulator

where `array_type` is a double-precision floating-point

```
1 #typedef double array_type;
```

With this code, the HLS compiler *synthesizes* a circuit that executes our instructions. What is remarkable is that we can implement this circuit with little to no knowledge of hardware design. In contrast, attempting to represent this simple example using the RTL approach proposed in Peri (2020) cannot abstract from these technicalities. To demonstrate this point, Listing 1 in Supplemental Appendix A.1 reports the RTL code required to implement the same circuit above.

Beyond the complexity of the syntax itself, coding in RTL is akin to describing a circuit. Accordingly, it demands proficiency in digital design, a subject rarely within the purview of economists. In contrast, the HLS approach delegates these details to the compiler. More significantly, the compiler autonomously recognizes and implements many

of the performance optimizations discussed in the next sections without explicit directives from the user.

4.1.1 FPGA optimizations Our previous section implemented an accumulator without any optimization. This approach is often referred to as the “dusty deck” approach, evoking the image of pulling an old deck of punched cards off the shelf and compiling it, untouched, to the new hardware.

While fully functional, the synthesized circuit is not efficient. Figure 1 illustrates the associated data flow graph to clarify this point. The vertical axis denotes the passage of time in clock cycles, while the horizontal axis represents the circuits’ functional units, consisting of additions and memory accesses in our example. The graph shows that every P clock cycles, a new double-precision floating-point addition is performed. Accordingly, the synthesized circuit has a for loop latency of $J \times P$ clock cycles: this is the time (measured in clock cycles) required for a circuit with a single physical adder to sequentially execute $J = 8$ additions (loop iterations) when each one of them takes $P > 1$ clock cycles (iteration latency).

We can enhance the efficiency of this circuit by reducing its for loop latency. The next sections introduce the main tools that economists can leverage in order to *assist* the HLS compiler to design an effective CL design. We use the word assist deliber-

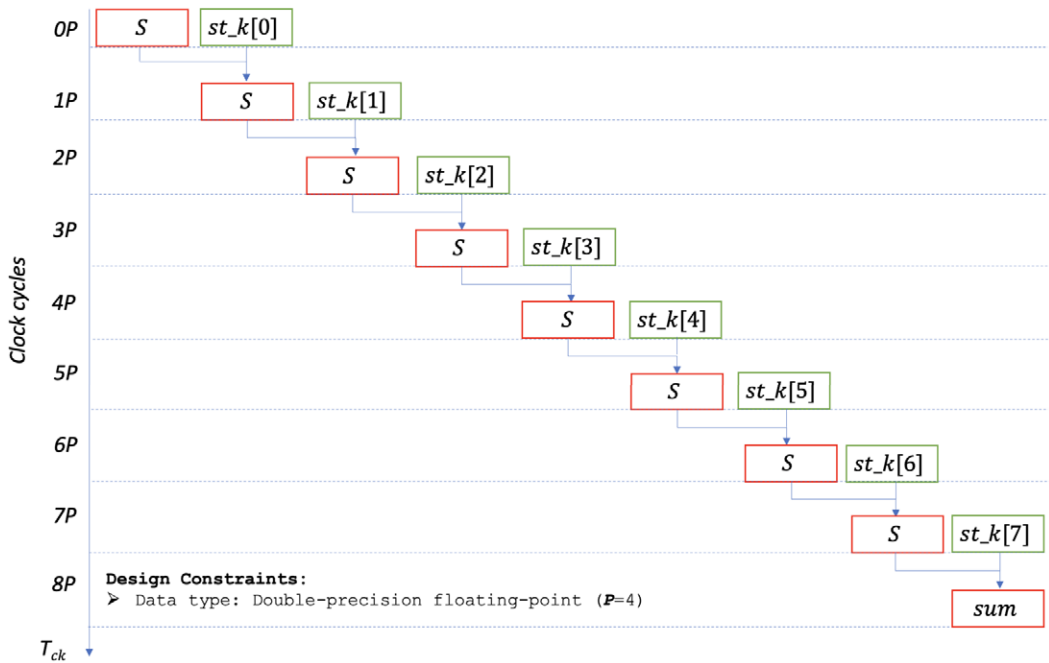


FIGURE 1. Data flow of sequential accumulator: Dusty deck approach. *Note:* Data flow graph of a sequential accumulator circuit for adding double-precision floating-point elements of an array st_k of size $J = 8$. The vertical dimension illustrates in which clock cycles these operations are performed (*scheduling*). The circuit consists of one single double-precision floating-point adder, which performs an addition every $P > 1$ clock cycles.

ately. Unlike in the RTL approach, we do not design the circuit. Simply, we provide clues about acceleration opportunities that the compiler may not autonomously identify.

We introduce these optimization techniques in an order that aligns with industry standards for hardware design optimization. Our overarching strategy is to initially establish an efficient circuit (Sections 4.2 and 4.3) and subsequently enhance its performance by harnessing parallel computing opportunities (Section 4.4).

Throughout this process, we address potential bottlenecks that might hinder efficient design. This enables us to compile a list of critical hardware issues that economists are likely to encounter when accelerating their applications. Most importantly, it provides the core hardware design principles for accelerating economic models.

4.2 Arbitrary-precision fixed point

We start our optimizations by targeting the `iteration` latency, P . This is the time required to execute the instructions in the body of the loop of Listing 1 at any given iteration. In our application, P consists of the time necessary to perform two operations: reading an array element and accumulating it (see Figure 1). Here, we illustrate how we can reduce the execution time of the accumulation from multiple clock cycles to a single one by transitioning from single or double-precision floating-point data type (hereafter referred to as floating point) to fixed-precision fixed point (hereafter referred to as fixed-point) to represent the array elements.

Floating-point operations provide the IEEE754 standard for the finite precision approximation of real numbers. However, they cost hardware and energy and take multiple clock cycles to complete. In contrast, integer and fixed-point operations execute within a single clock cycle and, when properly implemented, deliver high accuracy.⁸ Among others, Yates (2009) provides guidelines for employing fixed-point arithmetic, and Section 7.2 operationalizes them in our main application using HLS (Xilinx, Inc. (2021)).

Here, we illustrate how a few lines of code can recast floating- into fixed-point arrays by defining a type `fixed_t` with the required precision in a header file, along with the necessary libraries.⁹

```

1 #define FIXED_ACC 1
2 typedef double real;
3 #if FIXED_ACC
4     #include <ap_int.h>
5     #include <ap_fixed.h>
6     typedef ap_fixed<72, 21> array_type;
7 #else
8     typedef real array_type;
9 #endif

```

Conversely, the implementation of fixed-point arrays in RTL can be tedious even for experienced hardware designers. For example, the RTL designer must explicitly manage

⁸See Supplemental Appendix A.2 for a brief overview of finite-precision fixed-point approximation.

⁹This is one of the few optimizations that the HLS compiler cannot perform automatically since it requires users to specify the fixed-precision parameters (72 and 21 in line 6). The macro `FIXED_ACC` allows flexible toggling between double-precision floating point and fixed point for debugging purposes.

the proper alignment and interpretation of values with different fixed-point precision representations.

Besides reducing the iteration latency, recasting our operation in fixed point also addresses two well-known bottlenecks that hinder the efficient design of an accumulation with addition operations: (i) memory carry dependency; and (ii) the lack of associativity property of floating-point additions. Section 4.3 demonstrates how fixed point addresses the memory carry dependency problem, facilitating the creation of efficient pipelines. Section 4.4 shows how properties of fixed-point arithmetic enable the implementation of a parallel reduce tree to solve our accumulation, further reducing the `for loop` latency.

4.3 *Pipelining and the quest for an initiation interval of 1*

An ideal CL design keeps the circuit constantly busy (*goal 1*) and starts new computations at every clock cycle (*goal 2*). Here, we present an optimization tool that helps achieve both goals: hardware pipelining. In this context, we discuss how our solution of recasting our accumulator from floating to fixed point tackles a key obstacle to the efficient pipelining of our accumulation circuit: the memory carry dependency problem.

4.3.1 Pipelining To understand hardware pipelining, we need to introduce the concept of *initiation interval*, denoted as Π . The Π measures the number of clock cycles elapsed between the launch of two consecutive iterations in a loop. An ideal circuit has an $\Pi = 1$, and *initiates* computations at every clock cycle. The following code listing uses the `#pragma HLS pipeline $\Pi = 1$` directive to synthesize this design.

```

1 void hw_loop(array_type st_k[J], array_type &reduced_sum){
2     array_type sum = 0.;
3     for(int i=0; i<J; i++) {
4         #pragma HLS pipeline  $\Pi = 1$ 
5         sum+=st_k[i];
6     }
7     reduced_sum = sum;
8     return;
9 }
```

LISTING 2. Pipelined accumulator

To understand the hardware implications of this optimization technique, Figure 2 reports the data flows associated with four distinct hardware designs of our sequential accumulator: floating point (Section 4.1), fixed point (Section 4.2), fixed point with $\Pi = 1$, and floating point with $\Pi = 5$. These graphs illustrate how the different CL designs engage their hardware components, often referred to as *pipeline stages*. They do so by reporting which data points each pipeline stage (horizontal axis) processes during each clock cycle (vertical axis).

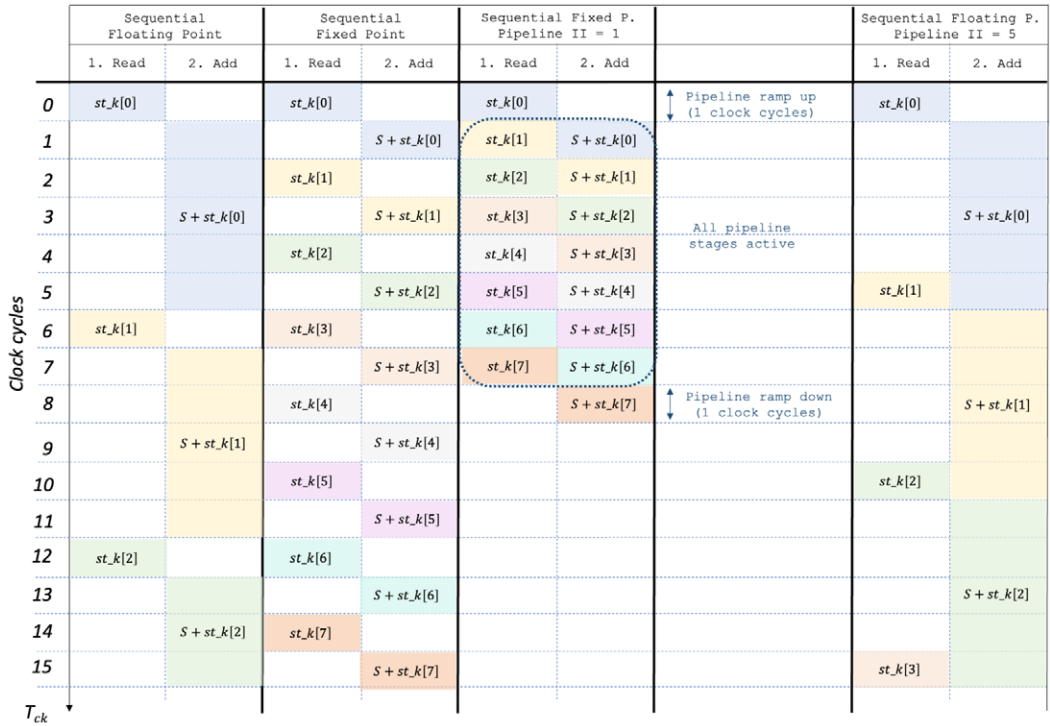


FIGURE 2. Data flow of sequential accumulator with pipelining. *Note:* Data flow of four different hardware designs, from left to right: sequential floating point (double precision), sequential fixed point, sequential fixed-point pipeline with $\text{II} = 1$, sequential floating-point (double precision) pipeline with $\text{II} = 5$.

In the case of our sequential accumulator, the pipeline stages associated with the hardware implementation of the body loop's instructions (line 5 of Listing 2) consists of two operations: a reading and an accumulate operation. Later in our main application, these pipeline stages will consist of the hardware steps required to solve either the **IHP** problem in equation (8) or the **Simulation** step.

The first two data flows on the left side of the figure illustrate how transitioning from floating to fixed point significantly reduces the execution time of the adder, from five to a single clock cycle, and thus, the iteration latency from six to two clock cycles. Consequently, the total for loop latency drops from $P \cdot J = 6 \cdot 8 = 48$ clock cycles to $2 \cdot 8 = 16$ clock cycles.

Listing 2 improves over the sequential fixed-point design by instructing the HLS compiler to synthesize a fixed-point accumulator that initiates computations every clock cycle ($\text{II} = 1$). This command generates a circuit that operates as follows. At clock cycle 0, the circuit starts reading $\text{st_k}[0]$ from memory. At clock cycle 1, the circuit initiates the first fixed-point addition $S = S + \text{st_k}[0]$, while simultaneously loading $\text{st_k}[1]$ for a new accumulation. From clock cycle 1 to 7, the circuit reaches its maximum degree of *instruction level parallelism*. At this point, all functional units of the circuit (pipeline stages) are working in parallel (attaining *goal 1*), yielding an addition

at every subsequent clock cycle (attaining *goal 2*). The time required for all `pipeline` stages to become active is called `pipeline ramp up` and equals one clock cycle in our example. After clock cycle eight, the pipeline stages start to empty, as there are no more array elements to read. By clock cycle nine, all pipeline stages are empty. The time required to flush out the pipeline stages is called `pipeline ramp down`, also equal to one clock cycle in our application.

As the state space increases, the pipeline ramp-up and ramp-down times become negligible with respect to the total for-loop latency. For example, when transitioning from an accumulation of an array of size 8 to an array of size 10,000 in equation (12), the portion of `for loop latency` spent on filling and flushing the pipeline shrinks from 25%, $\#Stages_{\text{ramp up} + \text{ramp down}}/J = 2/8$, to a negligible 0.02% ($2/10,000$).

Memory carry dependency problem Add-reduce accumulations inherently suffer from memory-carry dependencies because adding a new element of the array to the accumulation requires the previous element of the array to have already been incorporated (Figure 1). In the context of floating-point accumulations, this data dependency may impose a constraint on the lag (in cycles) at which new computations can be initiated (**II**).

For example, attempting to design a circuit with an **II** = 1 would result in a synthesis failure, flagging a memory carry dependency violation. The reason becomes evident when inspecting the behavior of this circuit on the right-hand side of Figure 2. The circuit must wait until clock cycle 5 for the value $S + st_k[0]$ to become available before proceeding to the new addition $S + st_k[1]$. Consequently, the best initiation interval for this design is every five clock cycles (**II** = 5). Transitioning to fixed point effectively addresses the memory dependency bottleneck by reducing the waiting time for the next element of the accumulation from five to just a single clock cycle, reaching **II** = 1.

4.4 Loop unrolling and memory management

So far, we have used HLS to design a circuit that efficiently initiates the necessary computations for a sequential accumulator at every clock cycle. Next, we further enhance its performance by introducing a powerful HLS tool that allows us to harness parallel computing opportunities: loop unrolling. Alongside this, we address limitations associated with memory management and provide effective solutions. Furthermore, we discuss a common performance hurdle faced by the hardware implementation of an accumulator, such as equation (12): the lack of associativity in floating-point additions. We conclude by explaining how our solution of recasting our accumulator from floating to fixed point effectively overcomes this challenge.

4.4.1 Loop unrolling Our prior circuitry executed additions sequentially, overlooking the opportunity to exploit the data independence of array elements for concurrent parallel addition operations. In our accumulation example, and later in the acceleration of equation (8) and Simulation step of the Krusell and Smith (1998) algorithm, we harness this parallelism by deploying the `#pragma HLS unroll factor = X`. This directive creates multiple physical replicas (X) of the circuit described within the loop's body (line 6 of Listing 3) and *autonomously* organizes them to operate in parallel in a reduce-tree

pipeline, which we will describe momentarily. Intuitively, the more physical replicas operating in parallel (`unroll factor`), the fewer iterations are necessary to perform the final accumulation, lowering the `for loop latency`. The following listing synthesizes a pipeline that processes in parallel all elements of the array, $J = \text{unroll factor} = 8$.

```

1 void hw_loop(array_type st_k[J], array_type &reduced_sum){
2     #pragma HLS array_partition variable = st_k factor = 4 type = cyclic
3     array_type sum = 0.;
4     for (int i=0; i<J; i++) {
5         #pragma HLS unroll factor=8
6         #pragma HLS pipeline II = 1
7         sum+=st_k[i];
8     }
9     reduced_sum = sum;
10    return;
11 }

```

LISTING 3. Accumulator with loop unrolling of factor 8

Associative reduce tree and fixed-size loop bounds Figure 3 shows that an unroll factor of 8 prompts the compiler to *autonomously* generate an associative reduce tree structure with a pipeline depth of three layers, resulting in an accumulation every three clock cycles.¹⁰ This design is facilitated by the fact that the loop bounds are fixed (0 and J), and increments of the loop iterator i are constant (and equal to 1). These clues allow the compiler to aptly determine the required number of layers: $O(\ln_2 \text{unroll factor})$.

When the array of size J is larger than the `unroll factor`, this circuit produces an accumulation every $J/\text{unroll factor} + \ln_2(\text{unroll_factor}) + 1$ clock cycle. The first term illustrates how loop unrolling decreases the `for loop latency` from J to $J/\text{unroll factor}$ by placing enough hardware, as determined by `unroll factor`. The second term captures the pipeline ramp-up time.¹¹ As discussed in Section 4.3, the pipelining performance of this circuit increases with the size of the state space. As J increases, the pipeline ramp-up time becomes negligible with respect to the total accumulation time, which converges to $J/\text{unroll factor}$.

While the acceleration achieved with loop unrolling resembles the single instruction, multiple data acceleration exploited by GPUs, it differs in important dimensions. First, the HLS compiler designs new hardware with the requested level of parallelism,

¹⁰In the first clock cycle, the compiler simultaneously reads all elements of the array. In the second clock cycle, it performs four additions in parallel. In the third clock cycle, it performs two additions and reduces the accumulation to a final digest. The fourth clock cycle appears when $J > \text{unroll_factor}$, in which case the accumulation of 8 elements is added to the overall rolling sum. See Supplemental Appendix A.3 for more details.

¹¹The third term appears when $J > \text{unroll factor}$ and captures the time required to incorporate the result of the reduce tree into the rolling accumulation (clock cycle 4 in Figure 3). Otherwise, the term disappears.

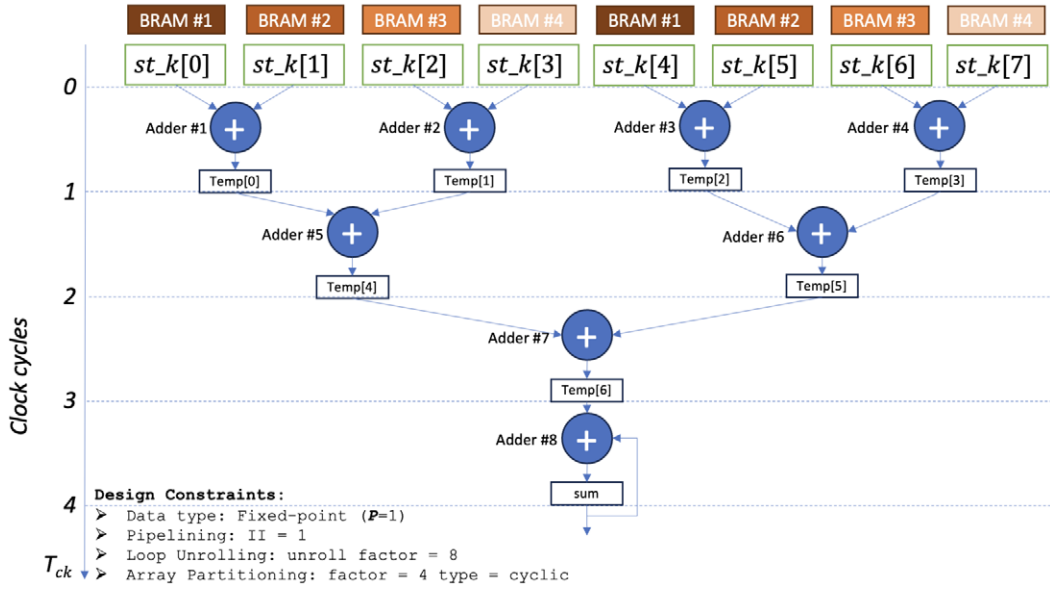


FIGURE 3. Data flow of accumulator with loop unroll. *Note:* Data flow graph of a circuit that unrolls by a factor of 8 the addition of the fixed-point elements of an array st_k of size $J = 8$. The vertical dimension illustrates in which clock cycles these operations are performed (*scheduling*). The circuit consists of eight fixed-point adders, which perform an addition every clock cycle. BRAM #1–#4 denote four local memories used to partition the st_k array.

whereas GPUs exploit existing one. Consequently, while GPUs may experience performance degradation when the required level of parallelism does not effectively engage all their cores (Aldrich et al. (2011)), FPGAs are affected to a much lesser extent. Second, FPGAs allow the implementation of hardware designs of intricate algorithms, whereby various facets of the application can exploit different acceleration strategies. The Krusell and Smith (1998) algorithm exemplifies this complexity, transitioning from a policy function iteration algorithm that benefits from reduced evaluation latency through instruction-level parallelism and pipelining (Section 4.3) to a multiparameter exploration or multiple-actor simulation step, which benefits from data-level parallelism. While GPUs thrive in the latter, they cannot reduce evaluation latency.

To conclude, note the `#pragma HLS array_partition` directive at line 2 of Listing 3. Next, we discuss how this command addresses a potential issue related to memory management that *may* prevent HLS from synthesizing the required circuit.

4.4.2 Memory management Now, we discuss two main limitations regarding memory management: large memory access latency and array partitioning.

Large memory access latency The FPGA device and host processor share a common memory referred to as global memory (dynamic random access memory, DRAM). Global memory is large (tens of gigabytes), and thus, useful for storing input, intermediate, and output data. In our accumulator design, we initialize our array in double-precision float-

ing point in the host processor, `st_in`, and make it available to the FPGAs by storing it in global memory. Access to this memory is slow (hundreds of clock cycles).

To solve this problem, we can copy chunks of data from the large global memories to on-chip local memories, which are small but numerous and can be accessed in a single clock cycle: block RAMs (BRAMs), ultraRAMs (URAM), LUT RAMs, registers. The following listing illustrates how straightforward it is to copy the array from global (`st_in`) to local memories (`st_k`) through the implementation of a simple array assignment.

```

1 void hw_top_init(real *k_in, array_type st_k[J]){
2     init_1:
3     for (int i = 0; i < J; ++i)
4         st_k[i] = (array_type) k_in[i];
5     return;
6 }
```

LISTING 4. Store data from global to local memories

Note the type cast operation that converts the floating-point values of the array initialized by the host, `st_in`, into the fixed-point array declared in the local memories, `st_k`. Without loss of generality, we assume that `st_k` has been stored in BRAMs in all our designs.

Array partitioning When transferring data from global to local memory, compilers optimize on-chip resource use by storing contiguous elements of `st_k` in the same BRAM before moving on to the next. For instance, without further instructions, Listing 4 stores all the elements of `st_in` in a single BRAM, which without loss of generality, we refer to as BRAM #1.

Each BRAM provides two read-write ports, allowing for a maximum of two concurrent read–write operations. Accordingly, while resource-efficient, this storage scheme can lead to a memory bottleneck when the total number of required reads from multiple pipelines (determined by the `unroll` factor multiplied by the simultaneous reads per pipeline) surpasses the available reading ports per shared BRAM.

The memory-access bottleneck becomes evident in our example when we instruct the compiler to synthesize four physical adders in parallel (Listing 3). In this scenario, during the first clock cycle, we can only read a maximum of two out of the eight `st_k` values stored in the single BRAM. Consequently, the synthesis process fails to produce the required CL design. Instead, it results in a circuit with an accumulation process that resembles the sequential addition described in Section 4.3 (see also Figure A.3 in Supplemental Appendix A.3.1).

The memory bottleneck can be resolved by strategically distributing the elements of `st_k` across multiple BRAMs to accommodate the eight memory accesses required by the four parallel pipelines. To achieve this, we use the `#pragma HLS array_partition` directive at line 2 of Listing 3. This directive employs a `cyclic` array

partitioning with a `factor` of 4 to interleave the elements of `st_k` across BRAM #1–#4, as depicted in Figure 3.¹²

Automatic optimizations We explain here the emphasis on the word *may* at the end of Section 4.4. In our simple accumulator example, for a given `unroll factor`, the HLS compiler recognized the need to partition the array `st_k` without explicit clues from the array partitioning `#pragma`. As compiler technology evolves, we anticipate such automation to expand to even more intricate circuits. Our performance analysis of the Krusell and Smith (1998) algorithm in Section 7 demonstrates the advanced nature of these automations already in place.

4.4.3 Nonassociativity of floating-point addition Here, we discuss how using floating-point additions prevents the compiler from autonomously generating the associative reduce tree displayed in Figure 3. Computers carry out computation on numbers with finite representations. When we abandon the field of real numbers $(\mathbb{R}, +, \times, \geq)$ for their IEEE754 floating-point, finite-precision approximation, we lose the associative property of the addition. Goldberg (1991) provides an example by showing how $(x + y) + z$ and $x + (y + z)$ give different results when $x = 1\text{e}30$, $y = -1\text{e}30$, and $z = 1$. The first equals 1, and the second equals 0 (since $y + z$ rounds to $-1\text{e}30$ before performing the addition with x). Accordingly, compilers have historically been prevented from rearranging floating-point additions (like in Figure 3) in order to avoid producing a different result than the original code. This rule downplays the fact that the source code may have been written in a loop simply for convenience, without a specific order of additions in mind. As a result, the compiler still insists on performing the floating-point sum sequentially in $J \cdot P$ cycles, as shown in Figure 1, even if we allowed it to use more hardware.

In our accumulator example (later deployed in the Krusell and Smith (1998), application), we tackle this issue by performing the accumulation in fixed point, where the associative property still holds, and the compiler knows it is safe to reassociate the additions since the computed result will be the same. For alternative strategies to handle this issue while maintaining floating-point, we refer to Hrica (2012) and Kadric, Gurniak, and DeHon (2016). In contrast to these strategies, our fixed-point approach (Section 4.2) also addresses the long latency of floating-point additions with minimal compromises in accuracy, as demonstrated by our application.

4.5 Portability, code reusability, and programming languages

C-code designed for FPGA acceleration using HLS exhibits remarkable portability. As our accumulation example demonstrates, the code compiles and runs on CPUs (including vector and SIMD computations) without modifications. Non-FPGA compilers ignore the

¹²Cyclical array partitioning creates `factor = 4` smaller arrays from `st_k` by copying elements cyclically, and storing them in independent BRAMs. For instance, `st_k[0]` goes into BRAM #1, `st_k[1]` into BRAM #2, `st_k[2]` into BRAM #3, `st_k[3]` into BRAM #4, `st_k[4]` into BRAM #1 and so on until completion. For a discussion of the different array partitioning types (cyclic, block, and complete), we refer to <https://docs.xilinx.com/r/en-US/ug1399-vitis-hls/pragma-HLS-array-partition>.

HLS `#pragmas`, eliminating the need to maintain a separate codebase for FPGA acceleration.

Furthermore, high-level computational components, such as accumulation (discussed here), interpolation, and simulation (discussed later) are common in economic models. Much like economists use function libraries to address recurring algorithms, the same HLS-enhanced codes and associated hardware components can be repurposed to accelerate recurring aspects of different models.

While HLS-enhanced C-code is highly portable, FPGA platform-specific optimizations and design may still require special considerations. First, fine-tuning an application for a specific FPGA involves adjusting parameters like loop unrolling or array partitioning, considering both application-specific and platform-specific factors, such as the execution frequency of particular portions of code and their resource utilization. These refinements only require changes to specific arguments in the `#pragmas`, leaving the base C-code unchanged. Recent advancements in automatic compilation and optimization hold great promise for automating these optimizations with solvers and autotuners (as discussed below), removing the need for manual changes by the developer.¹³ Second, interfacing with platform-specific components, like memory interfaces and I/O ports, super logic regions (SLRs) may be necessary for advanced designs (see Section 5.1.2). To mitigate these platform-specific concerns, we showcase FPGA acceleration on AWS, a cloud-based service that offers widespread access to *identical* high-performance FPGA devices. In addition, our tutorial provides *all* the necessary boilerplate code to harness these platform-specific components.

To conclude, recent years have seen great advances in the incorporation of hardware design tools among popular programming languages, like Python, Julia, and Matlab (e.g., Quenon and da Silva (2021), Biggs, McInerney, Kerrigan, and Constantinides (2022)). Our accumulator best exemplifies these efforts. The Matlab suite has incorporated an HDL coder that generates the RTL code of a variety of functionalities, including the accumulator we just described in this section.¹⁴ Overall, high-level synthesis compilers are bound to become more popular among high-level software languages, further easing economists' access to custom accelerators.

5. OUR APPLICATION

This section illustrates how we deploy the FPGAs' optimization tools introduced in our previous section to accelerate the Krusell and Smith (1998) algorithm in Section 3. We present the FPGA acceleration approach in Section 5.1 and the CPU counterpart in Section 5.2. Supplemental Appendix B reports the hardware specifications of the different AWS instances.

¹³We can determine these tuning parameters as the ones that maximize performance while staying within the constraints of available FPGA hardware resources. Solvers or autotuners can then automate the selection of these values, streamlining the optimization process (Whaley and Dongarra (1998), Kapre and DeHon (2009), Lo and Chow (2016), Zhao, Feng, Sinha, Zhang, Liang, and He (2020)).

¹⁴See <https://www.mathworks.com/help/fixedpoint/ref/sum.html>.

5.1 *Hardware architecture of FPGA acceleration*

The FPGA acceleration approach solves the [Krusell and Smith \(1998\)](#) algorithm as implemented by [Maliar, Maliar, and Valli \(2010\)](#) in Section 3 by using Amazon F1 instances to deploy three configurations of FPGAs connected to a host: one FPGA (f1.2xlarge), two FPGAs (f1.4xlarge), and eight FPGAs (f1.16xlarge).

The computational flow is as follows. The host initializes variables, allocates jobs across available FPGA(s), and transfers the data to the FPGA(s). The FPGA(s) then execute(s) the *kernel*: the nested, fixed-point algorithm detailed in Section 3. Subsequently, the host reads back and saves the results. Unless differently specified, computations employ the IEEE754 double-precision floating-point format.

5.1.1 *FPGA hardware description* F1 instances mount (up to eight) Xilinx VU9P FPGAs, each with 1.2 million 6-input gates (LUTs), 6840 multiply-accumulate units (DSPs), and 346Mb of on-chip memory, given by 76Mb of block RAM (BRAM) and 270Mb of ultra RAM (URAM). These resources are organized in three dies, referred to as a SLR, connected by a silicon interposer (see Figure 5). The AWS FPGA is further divided into two partitions: shell and custom logic (CL). The shell is the FPGA platform implementing external peripherals, such as peripheral component interconnect express (PCIe) and dynamic random access memory (DRAM). Users' CL designs can employ up to 895 thousand LUTs, 5640 DSPs, and 284 of on-chip memory, given by 59 Mb of BRAM and 225 Mb of URAM per FPGA.

5.1.2 *Custom logic hardware design* Now, we operationalize the FPGA optimization tools discussed in Section 4 to accelerate the [Krusell and Smith \(1998\)](#) algorithm. In addition, we show how our CL design exploits platform-specific characteristics of the F1 instances' FPGAs: the presence of three SLRs. As mentioned above, the Xilinx VU9P FPGAs are organized in three dies, whose custom logic can be jointly or independently used for acceleration purposes. For our application, we tailor our CL design to compute three economies in parallel, one per SLR, in the same FPGA.¹⁵

[Cheela et al. \(2023\)](#) describe in detail how the OpenCL commands organize the computation flow in kernels. Each kernel is assigned a different economy θ and is instantiated in a separate SLR. Kernels are then deployed in parallel. A feature of this hardware design is that it is easily scalable with minimal modifications of the C/C++ code. OpenCL commands collect the available SLRs and instantiate the kernels. In the case of one FPGA (f1.2xlarge), three kernels are instantiated across the available SLRs. In the case of eight FPGAs (f1.16xlarge), 24 kernels are launched in parallel.

To load and deploy our custom logic hardware design on the EC2 F1 instances, we create Amazon FPGA Images (AFI) that combine the shell and the CL design for multiple grid sizes. We accomplish this step by creating AWSXCLBIN files.

¹⁵We could alternatively design the CL to span across the multiple SLRs. This design would require handling inter-SLR connections, complicating the coding, and the management of tighter clock constraints. Since many of the benefits of accelerators are associated with the estimation of structural parameters that involve the computation of several thousand economies (each associated with a different set of parameter values), our hardware design is appropriate. The single SLR solution is still 28.38times faster than the sequential execution in the CPU (Section 6.5).

5.1.3 The single-kernel design Our kernel implements in hardware the nested fixed-point algorithm presented in Section 3. The design organizes the three interdependent building blocks—IHP, simulation, and ALM—in a sequence, and it performs the computations until convergence of the ALM coefficients in equation (13).

Following Amdahl's law, our CL design allocates hardware resources to balance the time spent on the two time-consuming steps of the algorithm: the **IHP** and Simulation design. Our final kernel design starts computations at every clock cycle at 250 MHz and provides flexibility to study multiple grid sizes.

To reach an Π of 1, we address the challenges discussed in Section 4 as follows. First, we tackle the two memory management bottlenecks: (i) large memory access latency; and (ii) limited memory ports to access data in parallel. We fix the first issue by copying data in local memories that can be accessed in a single clock cycle. Our FPGA has enough local memories to store all of the input data for our application (Table A.3). Hence, we need to transfer the input data only once per kernel computation. We address the second issue by deploying the array partitioning `#pragma` to store multiple copies of the same data in as many independent local memories as required by the number of pipelines performed in parallel.

Next, we turn to application-specific interventions essential for achieving an Π of 1, specifically the acceleration of the linear interpolation within the **IHP** and **Simulation** steps. A well-known computational challenge in interpolation is to find the interval of interpolation. We accelerate this step by implementing an efficiently pipelined (Π of 1) jump search algorithm with fixed-size loop bounds.¹⁶ To ensure that the algorithm does not introduce any unwanted bias, we verify that it outperforms alternative search algorithms in the CPU (Table 2).

The individual household problem (IHP) design The **IHP** design implements the previous steps to reach an Π of 1 in equation (8) by pipelining the operations involved in the computation of the expectation as discussed in Section 4.3. We further accelerate the design by using the partial loop unrolling technique discussed before to place two pipelines to work in parallel on the (sequential) solution of equation (8) at different states $\{k, \epsilon, m, A\}$ (more pipelines are hard given the hardware limits). Given the current guess of individual capital holdings, the IHP design solves equation (8) for every state and updates the guess according to equation (9), as discussed in Section 3.B.(ii). These steps are iterated until convergence, as per equation (10).

The simulation design The simulation design in Section 3.D(ii)(b) is divided into two steps, each accelerated with a custom pipeline. The first step is represented by the accumulation operation required to compute the cross-sectional average of the individual household's capital holdings, m_t , in equation (12). To achieve an $\Pi = 1$ at this step, we deploy the custom pipelined fixed-point accumulation operator discussed in Section 4. Indeed, the best acceleration performance for our benchmark model is reached by letting the compiler perform automatic optimizations. The second step is represented by

¹⁶As discussed in Section 4.4.1, the presence of *fixed* interpolation grids' bound allows the compiler to autonomously determine the depth of the reduce tree required to find the interpolation interval.

the interpolation step involved in the computation of next-period household capital holdings. We reach an $\Pi = 1$ by implementing the aforementioned jump search algorithm.¹⁷

The ALM design The ALM step involves several resource-expensive operations but has very small latency (1.25 ms in our baseline economy).¹⁸ Accordingly, we provide instructions to the compiler to not perform any pipeline or unrolling and further instruct it to limit the number of hardware resources used.

5.1.4 The three-kernel design Our CL design is tailored to compute three economies in parallel by exploiting the three hardware regions available for hardware design: the SLRs. To reach this goal, we modify our single-kernel design to handle the resource limitations. Our single-kernel design barely fits in one of the SLRs, slightly leaking into one of the adjacent SLRs with nontime-critical operations (Figure 4). The Amazon shell—which provides useful hardware resources to facilitate the HLS kernel design—exacerbates this problem by consuming part of the resources in two of the three adjacent SLRs. To fit the three kernels, we provide directives that limit the degree of partial unrolling in the IHP. The resulting three-kernel design (Figure 5) trades off the single-kernel performance for the parallel execution speedup.

5.2 Hardware architecture of CPU multicore acceleration

The CPU multicore acceleration approach solves our model using Amazon M5N instances by parallelizing data-independent tasks (economies) across multiple physical cores: one core (m5n.large), eight cores (m5n.4xlarge), and forty-eight cores (m5n.24xlarge). Supplemental Appendix B justifies the choice of these instances.

Our benchmark CPU kernel is an optimized C++-version of the original [Maliar, Maliar, and Valli \(2010\)](#) Matlab code, whose sequential single-core solution of the baseline economy is ten times as fast as the original Matlab code, an expected speedup given the results in [Aruoba and Fernández-Villaverde \(2015\)](#). The CPU kernel utilizes a jump search algorithm in order to determine the interpolation interval over the individual capital holdings grid. Table 2 shows that this algorithm outperforms standard alternatives in the CPU, ensuring the best CPU kernel performance.

We operationalize the multicore parallelization by using the open source message passing interface (Open MPI).¹⁹ Open MPI provides easily implementable, off-the-shelf routines for parallelizing data-independent tasks across multiple cores. In contrast

¹⁷To achieve an $\Pi = 1$ in both steps, we create multiple copies of the input data to avoid memory access bottlenecks. The simulation requires a large number of individual shocks ($10,000 \times 1100$). To minimize the memory usage, we encode the $\{0, 1\}$ shocks (from integers) into bits and pack them in groups of eight into a single byte of memory (8 bits) in both the FPGA and the CPU. Further, we tell the compiler to store these shocks in the URAMs, which have wide arrays of 72 bits and allow us to store 64 shocks (of size 1 bit) in each of these arrays. By doing so, we need to access the memory only once every 64 iterations.

¹⁸The ALM step requires 16 BRAM, 133 DSP, 11,494 LUTs, and 15,265 Registers.

¹⁹See <https://www.open-mpi.org>. We use Open MPI: Version 4.1.1.

TABLE 2. Benchmarking the CPU: alternative search algorithms.

	Linear Search	Binary Search	Jump Search
Solution Time	73,657.8	38,392.0	28,452.5
Speedup	–	1.92	2.59

Note: Solution time (in seconds) and speedups of alternative interpolation interval search algorithms. Speedups are computed relative to the linear search algorithm. Results are obtained by solving 1200 baseline economies sequentially using a single core instance (m5n.large).

with alternative multicore parallelization (like OpenMP), Open MPI does not require different cores to share the same memory. This feature makes Open MPI particularly appealing for massive data parallelization, from small clusters up to medium/large-scale supercomputers.

Our computational flow is as follows. Open MPI routines determine the number of cores available and uniformly spread the number of (data-independent) economies to be computed across the different cores. For example, the solution of 1200 economies on a 48-core machine would have Open MPI allocating 25 economies per core. Each core will then independently complete the assigned tasks.

6. QUANTITATIVE RESULTS

We assess the efficiency gains of FPGA acceleration against CPU cores by measuring efficiency gains in solution speedup, AWS costs, and energy savings. We choose as a benchmark the CPU rather than the graphic processing units (GPUs) for two reasons. First, CPUs are still more commonly used in economics than GPUs, making the interpretation of our results more transparent. Second, differently from the CPU, the effectiveness of GPU acceleration is heavily dependent on software optimizations (e.g., memory management), making the benchmark performance more susceptible to software engineers’ choices. Nonetheless, for completeness, we implement the Maliar, Maliar, and Valli (2010) algorithm in python/1.13 and accelerate it using the Numba Cuda compiler on an NVIDIA A100 GPU. Consistent with the finding in Aldrich et al. (2011) for GPU acceleration on small grids, the GPU is “only” four times as fast as the original Matlab code. While the implementation of the algorithm with C-cuda and additional memory optimizations might yield further speedups, it is most unlikely that GPUs can deliver the same performance as the FPGA for this particular application.

To make our acceleration comparison as meaningful as possible, we proceed as follows. First, our FPGA and CPU kernel share the same C/C++ code to solve the same algorithm in both platforms. Hence, the observed speedup is the result of a margin available to the electrical engineer economist (designing hardware) and not to the software engineer economist (writing better code for CPUs and GPUs). Using the optimization techniques discussed in Section 4, we program the logic of our FPGA to accelerate the Krusell and Smith (1998) algorithm. In comparison, the CPU and GPU logic cannot be programmed. The CPU is premanufactured to optimize the performance of sequential

instructions needed in daily computer activities, which include not only simulating our models but also browsing the internet, checking emails, and so forth. As such, it lacks the gains from hardware specialization that are accessible to FPGAs. That said, the C++ compiler on a CPU and GPU goes a long way to execute the code as fast as possible by *automatically* performing low-level instruction parallelism. We can instruct the compiler to do this as much as possible by compiling the code using the aggressive `-O3` optimization flag in our G++ compiler.

Second, our baseline comparison is that of a single FPGA chip against a single CPU core while computing N_E economies (which we will define momentarily). This benchmark provides an apple-to-apple comparison of the differential performance of the two accelerators.

Third, we introduce a multieconomy parallelism to get insights into the potential of FPGAs to structurally estimate our model by solving it for many different parameter values. But we do so in a controlled way. A well-known drawback of `Open MPI` parallelization is that the total execution time is determined by the execution time of the slowest core. To temper this effect, first, we pick a number of economies ($N_E = 1200$) that is uniformly divisible across the different CPU cores in our AWS instances (1, 8, 48 cores). Second, we determine the benchmark speedup (Table 3) by solving the same economy (i.e., the same set of parameter values) multiple times. In this way, we eliminate heterogeneity in the solution time attributed to differences in convergence of the iterative algorithm with different parameter values. Under this setup, we ensure that we use all CPU cores and that the speedup scales linearly with the number of cores. Incidentally, this exercise illustrates how easy it is to implement increasingly aggressive acceleration by deploying multiple FPGAs in parallel, a dimension not explored in Peri (2020).

In summary, our benchmarking strategy eliminates all differences in performance between FPGAs and CPUs that are not inherently linked to the differences in their architecture. If we could rewrite the solution algorithm more efficiently, that better code would improve the performance of both FPGAs and CPUs and leave the relative performance (roughly) unchanged.

Table 3 reports how the relative performance of FPGAs versus CPU cores varies as we vary the number of CPU cores and FPGA devices. We solve for 1200 times the baseline economy on AWS instances connected to one (f1.2xlarge), two (f1.4xlarge), and eight (f1.16xlarge) FPGAs. We compare these results with the ones obtained by solving the model on AWS instances with one (m5n.large), eight (m5n.4xlarge), and forty-eight (m5n.24xlarge) cores. Then we measure the relative performance across speedups, cost savings, and energy savings. To make the comparison as meaningful as possible, we compute the speedup using the solution time—which is the time required to solve the algorithm on the FPGA/CPU. This approach abstracts from the time absorbed by initializations and host-FPGA communications. Supplemental Appendix C.2 details the reasons behind this choice by also reporting the execution time, that is, accounting for nonkernel operations. Table A.4 in the Supplemental Appendix complements this information by reporting costs and energy consumption by instance.

TABLE 3. Efficiency gains and implementation costs of FPGA acceleration.

Panel A: Efficiency Gains of FPGA Acceleration									
CPU-Cores	Speedup			Relative Costs (%)			Energy (%)		
	FPGAs			FPGAs			FPGAs		
	1	2	8	1	2	8	1	2	8
1	68.54	137.09	548.56	20.23	20.23	20.22	6.02	6.02	6.02
8	8.80	17.61	70.46	19.69	19.69	19.68	5.86	5.86	5.85
48	1.48	2.96	11.83	19.55	19.55	19.54	5.82	5.82	5.81

Panel B: Implementation Costs of FPGA Acceleration				
Extra Lines of Code	Kernel		Nonkernel	
	Number	Percent (%)	Number	Percent (%)
	75	5.37	128	51

Note: Panel A reports speedups provided by the FPGA and cost and energy usage of the FPGA relative to the CPU. The results are obtained by solving 1200 baseline economies using AWS instances connected to 1, 2, and 8 FPGAs and using open-MPI parallelization on AWS instances with 1, 8, and 48 cores (rows). Speedup is obtained by dividing the total solution time in the CPU by that in the FPGA. Relative costs and energy are calculated using on-demand AWS prices and total energy consumption, and reported as FPGA usage as a percent of CPU usage. Table A.4 in Supplemental Appendix C reports the details. Panel B estimates implementation costs for both kernel and nonkernel segments of our codebase by reporting the extra lines of code required by the HLS-enhanced C code when compared to standard C code designed to be executed on the CPU using Open MPI.

6.1 Speedups of FPGA acceleration

As mentioned before, our baseline comparison is a single FPGA versus a single CPU core sequential solution. The FPGA acceleration delivers a speedup of 69 times. The computation time drops from 8 hours in the CPU to 7 minutes in the FPGA (Table A.4).

Next, we document the FPGAs’ performance against that of multicore CPUs. As many researchers have access to high-end laptops with at least eight cores, we first compare one FPGA with the 8-core acceleration. The computation time is reduced from approximately 1 hour in the CPUs to 7 minutes in the FPGA (Table A.4). FPGA acceleration also compares favorably with respect to the CPU multicore parallelization on the AWS instance with the largest number of cores (48): a single FPGA device is 1.48 times faster than an Intel Xeon (Cascade Lake, second generation) platform running 48 cores in parallel.

Scaling these acceleration gains is easy, as it requires minimal modifications of the code to deploy multiple FPGAs in parallel. Panel A of Table 3 shows that eight FPGAs in parallel reduce the solution time by 549x, 70x, and 12x when compared with the 1-core, 8-core, and 48-core acceleration, respectively.

Panel A of Table 3 shows how the differential performance of FPGAs versus CPU cores scales linearly in the number of FPGA devices and CPU cores. This result corroborates the effectiveness of our benchmarking strategy in eliminating all contamination due to parallelization.

6.2 Cost savings of FPGA acceleration

We now turn to cost savings. The AWS cloud pricing schedules provide market-based measures of the cost of solving an application across different hardware architectures. We compute these costs by multiplying the total solution time for the AWS instance on-demand prices (Table A.2). Panel A of Table 3 shows that our application solves in the FPGA instances at less than a quarter of the cost of the CPU instances. This result is relevant because the structural estimation of parameters may require computing up to one million different economies. Hence, moving from CPU to FPGA computing would reduce the estimation costs by hundreds of dollars (from \$784 to \$159). Albeit the rental cost of FPGA instances per hour is higher, the acceleration we document more than justifies the expense.²⁰

Also, note that the linear speedup performance, combined with Amazon AWS linear price schedules, yields approximately constant costs and energy efficiency across different combinations of FPGA devices and CPU cores.

6.3 Energy savings of FPGA acceleration

We determine the energy consumption (joules) by multiplying the total solution time for the power consumption (watts) of FPGA and CPU chips. We use the AFI management tool to measure the FPGA average power consumption: 33 watts per FPGA device. We use the procedure discussed in Supplemental Appendix C.2.1 to measure the CPU power consumption: 8 watts per CPU core. The energy consumed by FPGA chips is less than 6% of the energy consumed by CPUs.²¹ This statistic is relevant for organizations with in-house computational clusters (such as research departments at central banks), whose computational needs are often constrained by the power limits on the cluster installation. Moving from CPU to FPGA computing enables more computations with the same energy.

Furthermore, the back-of-the-envelope calculations in Supplemental Appendix D suggest that the associated energy savings may also be relevant to reducing the carbon footprint impact of research computing. As we describe in Supplemental Appendix D, it has been estimated that the RMACC Summit and Blanca Supercomputers at the CU Boulder Research Computing Center emit 838.78 metric tons of CO₂ in the atmosphere every year to provide 150 million core hours to their research community. This is equivalent to the CO₂ emissions of 168 cars per year. If (a big if) we assume a type of acceleration similar to the one measured in our experiment and the FPGA power consumption recorded on Amazon Xilinx VU9P, transitioning all of these CPU-intensive computations to FPGA chips would reduce the CO₂ impact of these major supercomputing facilities to 31.07 metric tons of CO₂, or sixcars per year.

²⁰Another cost comparison is with an in-house cluster. Even without entering into their purchase cost, clusters are expensive to maintain. In our application, a single FPGA chip can perform the same task as a cluster with 69 cores. This is a medium-to-high scale cluster whose maintenance requires an HPC specialist, with a salary averaging around \$89,000 per year in the US circa 2024; see Zip Recruiter (<https://www.ziprecruiter.com/Salaries/HPC-System-Administrator-Salary>).

²¹Since CPU power consumption is proportional to the number of CPU cores, $\text{Power}(\text{cores}) = P \cdot \text{cores}$, and the solution time is inversely proportional to the number of CPU cores, $\text{Solution Time}(\text{cores}) = T/\text{cores}$, the energy, up to first order, is constant across CPU cores, $\text{Energy}(\text{cores}) = P \cdot \text{cores} \cdot T/\text{cores} = PT$.

6.4 Implementation costs of FPGA acceleration

Panel B of Table 3 reports the number of lines required to implement the kernel and nonkernel portion of our application for FPGAs and CPUs. The kernel portion of the HLS-enhanced C-code requires 5.37% of additional lines with respect to C-code written to be executed on the CPU. With 75 extra lines of code (`#pragmas`), the user can achieve a 69x speedup with respect to the single-core CPU. The nonkernel portion of the code requires 128 more lines (51%) to implement the `OpenCL` communications between host and device. However, this code barely varies across applications and is available on our Github repository.

6.5 Robustness

Table 4 reports a battery of robustness tests to study the performance of our hardware design. Different from the `RTL` approach proposed in Peri (2020), the compiler approach allows us to implement these experiments with minimal modifications of the C/C++ code and hardware design configuration file.

In our first exercise (Panel A), we solve the baseline economy using the single-kernel design in one FPGA and compare its performance with that of the single-core sequential solution. This result suggests that our single-kernel design can accelerate model experimentation, a valuable feature in the early stages of a research project when the ingredients of the final model are not yet determined. The sequential computation of the 1200 economies using the single-kernel design is relatively faster than that of the three-kernel design. Theoretically, instantiating three single-kernel designs in the available SLRs should bestow a speedup of $28.38x \cdot 3 \text{ SLR} \approx 85x$, yet the three-kernel design “only” reaches a speedup of 69x (Table 3). This loss in performance is due to the lower

TABLE 4. Single-kernel FPGA versus single CPU core.

Panel A: Benchmark Model, $\{N_k, N_M\} = \{100, 4\}$						
FPGA-Time (sec)	CPU-Time (sec)	Speedup (x)		Relative Costs (%)		Energy (%)
0.84	23.71	28.38		48.86		7.49
Panel B: Speedup Across Grid Sizes						
Aggregate Capital, N_M	4			8		
Individual Capital, N_k	100	200	300	100	200	300
Speedup (x)	28.38	34.41	34.31	27.81	31.05	32.06
Relative Costs (%)	48.86	40.30	40.41	49.85	44.65	43.26
Energy (%)	7.49	6.18	6.19	7.64	6.84	6.63

Note: Figures are obtained by comparing the solution of 1200 economies using AWS instances connected to 1 FPGA and sequential CPU execution on a single core. Panel A focuses on the benchmark economy, $\{N_k, N_M\} = \{100, 4\}$. Columns 1–2 detail the average solution time (in seconds) to compute the benchmark economy in a single-kernel, single-device FPGA (f1.2xlarge), and a single-core instance (m5n.large), respectively. Columns 3–5 display the efficiency gains of FPGA acceleration in terms of speedup, costs (in percent), and energy savings (in percent), computed as described in Table 3. The FPGA average power consumption on a single-kernel design is 17watts. Panel B studies how speedup, relative costs, and energy consumption vary with the size (columns) of the individual household capital holdings grid (N_k) and aggregate capital grid (N_M).

unrolling of the **IHP** design performed to save resources to fit the three kernels. The lower speedup is associated with higher costs than the costs of the three-kernel design, at still half of the cost. Energy savings remain significant.

In our second exercise (Panel B), we study how the speed gains depend on the size of the grids. The panel illustrates speedup, cost, and energy savings for increasingly finer grids on individual capital holdings $N_k = \{100, 200, 300\}$ and aggregate capital $N_M = \{4, 8\}$. To ensure a clear comparison, we keep contrasting the performance of the single-kernel FPGA design with the sequential single-core CPU execution. The speedup improves as the grid size of individual capital holdings gets larger, with stronger benefits when transitioning from 100 to 200 grid points, which then plateau at 300 grid points. The performance when moving from 4 to 8 grid points is slightly lower but comparable. Crucially, these results are obtained by keeping the *same* hardware design—pipelining, array partitioning, and loop unrolling—as the benchmark model with $N_k = 100$ and $N_M = 4$. This analysis suggests the potential for tailored optimizations for larger grid sizes, currently disabled in our design.

7. INSPECTING THE MECHANISM

What explains the observed speedup of FPGA versus CPU multicore acceleration? We address this question by discussing the performance benefits of gradual modifications of our code controlling the three critical acceleration channels: pipelining, data parallelism, and recasting data from floating to fixed point. Since the bottlenecks faced in our design are common to most dynamic programming problems, our acceleration strategy provides easily transferable tools to accelerate a vast class of economic models.

We start by illustrating the performance of a baseline model, whose hardware image is created by automatic optimization of the HLS compiler (*Baseline Model*). Next, we build up the acceleration by resolving problems that prevent us from achieving an efficiently pipelined loop (*Pipelining Channel*). Once we achieve a pipelined kernel, we exploit available resources to instantiate multiple copies of the pipelined loops (*Within-Economy Data Parallelism Channel*). We conclude by instantiating the three-kernel design across available SLRs and FPGA devices to run multiple economies (*Across-Economies Data Parallelism Channel*).

7.1 Baseline

Column 1 in Table 5 reports the speedup of solving the model using the single-kernel FPGA baseline design and the single-core CPU sequential solution. The baseline design differs from the single-kernel design discussed above because it does not explicitly call for any user-defined hardware optimizations (pipelining, unrolling, data precision). The only optimizations present in this design are the ones automatically performed by the HLS compiler. The compiler indeed tries to optimize the memory layout (to reduce the memory bottlenecks) and the loop pipelining (by trying to unroll inner loops). To reduce the latter effect, we use directives to limit the amount of automatic unrolling.

TABLE 5. Speedup gains: acceleration channels accounting.

	Baseline	Pipelining	Data Parallelism	
			Within Economy	Across Economies
<u>Single-core execution</u> FPGA Solution	0.21	6.94	28.38	68.54
<i>CL resources utilization (%)</i>				
BRAM	6.01	7.14	21.31	44.29
DSP	7.75	9.68	31.13	55.32
Registers	3.99	5.12	12.00	25.71
LUT	5.96	9.20	25.21	57.03
URAM	5.50	5.50	5.38	16.50

Note: Column 1 reports the speedup for a kernel design where all acceleration channels are switched off (baseline). Columns 2–4 report the speedup associated with implementing efficient pipelines (column 2), introducing data parallelism in the kernel design (column 3), and instantiating three kernels in the same FPGA (column 4). The speedup (row 1) is computed by dividing the total solution time in the one-core CPU by the solution time in the FPGA. The acceleration in columns 1–3 is performed using a single-kernel, single-device FPGA (f1.2xlarge), where column 4 coincides with the single-kernel design. The acceleration in column 4 is performed by deploying the three-kernel design in parallel across the three SLRs in a single FPGA (f1.2xlarge). Averages are computed over 1200 economies, except for the Baseline and Pipeline designs, which for cost considerations are computed over 120 economies. Resources are measured (using Xilinx Vivado) as a percentage of the Xilinx VU9P FPGA's resources utilized by AWS images associated with the different designs (columns). *Available Resources:* BRAM (1680), DSP (5640), Registers (1,790,400), LUTs (895 thousand), URAM (800).

The FPGA solution is reached in 113 seconds, approximately fivetimes slower than the CPU solution, which is reached in 24 seconds. This result is not surprising. FPGAs operate at a slower frequency than CPUs (in our case, 250 MHz vs. 3.5 GHz). Absent user-defined interventions, the CPU should be faster. That said, the compiler goes a long way in optimizing the design, as we would have expected the FPGA to be 14x (3.5 GHz/250 MHz) slower than the CPU, if only due to differences in clock frequency.

This acceleration illustrates how FPGAs' gains are a by-product of hardware specialization and not the result of the chip being intrinsically faster. Slower but better-organized tasks in the FPGA deliver higher performance than faster but “poorly” organized execution of the same tasks in the CPU. The next subsections discuss how we deploy the optimization tools discussed in Section 4 to efficiently organize these tasks in assembly lines, and how we improve performance by placing multiple assembly lines in parallel.

7.2 Pipelining

Next, we discuss the main changes we made in order to pipeline the **IHP** efficiently and **Simulation** steps in the single-kernel design.

Interpolation We accelerate interpolation as follows. First, we declare the loop bounds of the individual and aggregate capital grids (namely, $\{0, N_k\}$ and $\{0, N_M\}$) as fixed constants, allowing the compiler to autonomously physically *place* the required CL resources. Next, we implement a jump search algorithm to find the interpolation interval over the individual capital grid. The compiler instructs the hardware to pipeline a parallel reduce tree algorithm with three stages. Each stage determines the index of the

smallest grid value larger than the interpolation point $k'(k, \epsilon, m, A)$ by performing comparisons in parallel. The number of comparisons varies by stage and grid size and ensures that the entire grid is examined, $i = \{0, \dots, N_k\}$. The winner of each stage determines the search area of the successive stage. Since the result of this operation is part of a pipeline where the only dependence on subsequent loop iterations is through a final accumulation, we achieve an **II** of 1 in the **IHP** step.²²

Accumulation data precision The efficient design of the interpolation function accelerates both the **IHP** and **Simulation** steps. While the **IHP** step attains **II** = 1, the **Simulation** step has an iteration-to-iteration limit in the floating-point computation of the cross-sectional average of individual capital holdings, m_t , and settles at an **II** of 5. Section 4 demonstrates how recasting the pipelined accumulator in fixed point addresses this issue, yielding an **II** of 1 in equation (12), and consequently in the **Simulation** step. Table 5 records a speedup of 6.94 with respect to a single-core CPU. Importantly, currently available CPUs do not provide access to user-defined fixed-point arithmetic.

Our acceleration strategy trades off the accuracy of results for speed. The precision of the accumulations that occur in the algorithm can tolerate fixed-point calculations for the following reasons. First, the inputs to the accumulation are all positive values such that there will not be cancellation, which can often degrade the precision when moving from floating-point to fixed-point calculations. Second, since these are accumulations, we expect the accumulator not to hold a small value but rather to converge to a value in a limited range and magnitude. In our baseline economy, the accumulator sums up to values between 15.371273672208304 and 404851.76387144416. Knowing the accumulator's range, we can determine the required precision. In particular, we need at least $D_1 = 6$ and $D_2 = 15$ decimal digits above and below the decimal point to represent 404,851.76387144416 and 15.371273672208304, respectively. That is, we need $\lfloor \log_2 10^{D_1} \rfloor + 1 = 20$ binary digits to represent 404,851, and $\lfloor \log_2 10^{D_2} \rfloor + 1 = 50$ binary digits to represent 0.371273672208304. Accordingly, the minimum number of digits to represent our accumulator range is $50 + 20 = 70$. In order to accommodate a larger range of values (as may be required when we change economies θ), we set the number of digits to 72.

Not surprisingly, Table A.6 in Supplemental Appendix C.4 shows that the results are very accurate. The estimated ALM coefficients are identical up to the ninth decimal place (R^2 are all 0.999, and thus, not reported). The approximations of policy functions and distribution of individual capital holdings at $T = 1100$ are very good, as measured by the mean and max relative difference in the percent of these objects under floating and fixed point (less than $3.0e-08$). Moments of the distribution of individual capital holdings are identical up to the seventh decimal place. The Euler equation errors reveal that transitioning from computations in floating point on the CPU to fixed point on the FPGA has no discernible impact on the accuracy of our policy function's discretization (with relative differences lower than $1e-6$). The maximum Euler equation errors indicate that the accuracy of our solution increases with the individual capital holdings grid size, where the benefits of FPGA acceleration are even more pronounced.

²²In the CPU, the C++ compiler can autonomously decide to perform these operations in parallel, but the coder does not control this step.

7.3 Within-data parallelism

The interpolation and accumulation designs yield efficient custom pipelines but leave a considerable amount of CL resources unused in the single SLR (Table 5). Hence, our next step is to identify parts in the algorithm to parallelize using the loop unrolling technique described in Section 4.4.1. The computations involved in the interpolation step (Section 3.D(ii)(b)) of the Simulation and the policy function iteration (equation (8)) provide suitable candidates.

Given Amdahl's law, we perform a trade-off analysis between resource utilization and the solution speedup. Eight copies of the pipeline work well for our design, as they bring the solution time of the **Simulation** step (371 ms) closer to that of the **IHP** step (381 ms) when we unroll the latter over the states by a factor of two. Lacking hardware programmability, the CPU execution is more imbalanced, with 72% of the solution time spent on the **Simulation** step and only 28% on the **IHP** step. Accordingly, our CL design yields a 46x speedup in the **Simulation** step and 17x speedup in the **IHP** step compared to a single-core CPU, totaling a speedup of 28.38x as shown in Table 5. We call the design that follows from these changes the single-kernel design.²³

Also, it is clear from our acceleration strategy that the speedups at different stages of the algorithm are *endogenous* when programming FPGAs, where the hardware is not taken as given. Thus, FPGA acceleration depends less on the type of algorithms used compared to CPUs and GPUs. For example, replacing the stochastic simulation algorithm with a faster and more robust algorithm (e.g., Young (2010)) would allow us to free resources from this time-consuming step and reallocate them to accelerate the individual households' problem.

7.4 Across-economies data parallelism

We tailored our CL design to compute three economies in parallel. However, the single-kernel design consumes more resources than the ones available in the largest SLR, leaking into the adjacent SLRs with nontime-critical operations (blue area in Figure 4). Also, the AWS shell covers a good share of resources in the adjacent SLRs (orange area in Figure 4).

So, to fit three kernels, we slightly modify our design by reducing the usage of CL resources at the expense of performance. To do so, we provide directives to the compiler to halve the amount of unrolling. In particular, we reduce the unrolling in the **IHP** design from 2 pipelines working in parallel to a single pipeline. Figure 5 shows the usage of resources associated with the three-kernel design. Hence, we replicate the three-kernel design in each SLR and use OpenCL commands to launch one independent kernel in each of the three SLRs within a single FPGA device (f1.2xlarge). The FPGA design becomes 68.54times faster than a single-core CPU. Table 3 shows how we exploit this parallelism further to deploy more than one FPGA device in parallel on the f1.4xlarge and f1.16xlarge instances.

²³The recorded performance represents a lower bound to the performance that could be achieved by optimizing the individual grid size designs.

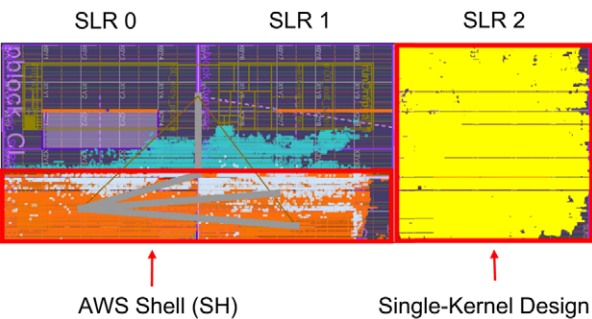


FIGURE 4. Single-kernel design: Resource utilization. *Note:* Resources utilized by: (i) the single-kernel CL design (yellow area); (ii) by the AWS shell (orange area); and (iii) available CL resources (other colors). The image is captured using Xilinx Vivado.

8. TOWARD ELECTRICAL ENGINEERING ECONOMICS

This paper proposes the design of FPGAs for the solution of economic models. This approach requires minimal knowledge of hardware design principles. With small modifications of standard C/C++ code, a single FPGA can deliver the same performance as 69 CPU cores when solving a canonical heterogeneous agent model. The associated energy savings make the compiler approach particularly appealing for organizations with in-house clusters whose computational needs are often constrained by the power limits on the cluster installation and the need to reduce the carbon footprint.

Our analysis leaves important venues for exploration. First, the recent popularization of machine learning techniques for the solution of economic models (Fernández-Villaverde, Hurtado, and Nuño (2019), and Kahou et al. (2021), among many others) may benefit from decades of research in the electrical engineering literature (Nurvitadhi et al. (2017)) in terms of the design of efficient FPGAs. A similar argument applies to the acceleration of maximum likelihood estimators.

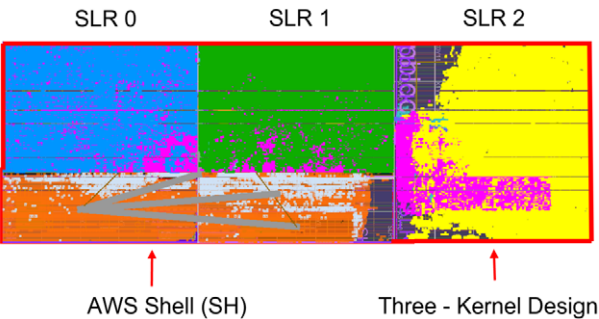


FIGURE 5. Three-kernel design: Resource utilization. *Note:* Resources utilized by: (i) the three-kernel CL design (yellow, green, blue areas each corresponding to one kernel); (ii) by the AWS shell (orange area); and (iii) available CL resources (other colors, of which the pink area serves as a wrapper). The image is created using Xilinx Vivado.

Second, notice that despite their acceleration potentials, FPGAs (like GPUs) still represent off-the-shelf application-specific integrated circuits (ASICs). Their routing network, logical units, and memory are predesigned by the manufacturer to be configurable, but they are not customized to serve any particular algorithm. With a growing literature and the development of sophisticated heterogeneous agent models to assess the effect of monetary policy or the economic impact of climate change, we foresee a not-too-distant future where central banks and other policymaking institutions would invest in the design of ASICs specialized in the solution of these models. It is hard to give a precise estimate, but customized silicons could likely improve the current speedup up to three orders of magnitude, with one order of magnitude only due to the faster clock cycle. FPGAs represent a first step in this direction, as they are actively used in the industry to test the functionality of the hardware design of customized chips. Of course, designing and manufacturing these pieces of silicon is not cheap (on the order of tens to hundreds of millions of dollars). Yet, the beneficial effects of a better-informed monetary or climate change policy dwarf these costs.

In conclusion, there is space for a new field, electrical engineering economics, focused on the design of computational accelerators for economics. Our analysis and successful experience in other areas suggest that such a field can provide computational breakthroughs in the years to come.

REFERENCES

- Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll (2021), “Income and wealth distribution in macroeconomics: A continuous-time approach.” *Review of Economic Studies*, 89 (1), 45–86. [0051]
- Aldrich, Eric M., Jesus Fernández-Villaverde, A. Ronald Gallant, and Juan F. Rubio-Ramírez (2011), “Tapping the supercomputer under your desk: Solving dynamic equilibrium models with graphics processors.” *Journal of Economic Dynamics and Control*, 35 (3), 386–393. [0050, 0066, 0073]
- Algan, Yann, Olivier Allais, and Wouter J. Den Haan (2008), “Solving heterogeneous-agent models with parameterized cross-sectional distributions.” *Journal of Economic Dynamics and Control*, 32 (3), 875–908. [0051]
- Algan, Yann, Olivier Allais, Wouter J. Den Haan, and Pontus Rendahl (2014), “Solving and simulating models with heterogeneous agents and aggregate uncertainty.” In *Handbook of Computational Economics*, Vol. 3, 277–324, Elsevier. [0051]
- Amman, Hans M., David A. Kendrick, John Rust, Leigh Tesfatsion, Kenneth Judd, Karl Schmedders, Cars Hommes, and Blake LeBaron (2018), *Handbook of Computational Economics*. Elsevier. [0050]
- Aruoba, S. Borağan and Jesús Fernández-Villaverde (2015), “A comparison of programming languages in macroeconomics.” *Journal of Economic Dynamics and Control*, 58, 265–273. [0052, 0072]

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2020), “Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model.” Working Paper 26647, National Bureau of Economic Research. [0051]

Azizi, Navid, Ian Kuon, Aaron Egier, Ahmad Darabiha, and Paul Chow (2004), “Reconfigurable molecular dynamics simulator.” In *12th Annual IEEE Symposium on Field-Programmable Custom Computing Machines*, 197–206. [0050]

Babb, Jonathan, Martin Rinard, Csaba Moritz, Walter Lee, Matthew Frank, Rajeev Barua, and Saman Amarasinghe (1999), “Parallelizing applications into silicon.” In *Seventh Annual IEEE Symposium on Field-Programmable Custom Computing Machines (Cat. No. PR00375)*, 70–80. [0051]

Bayer, Christian and Ralph Luetticke (2018), “Solving heterogeneous agent models in discrete time with many idiosyncratic states by perturbation methods.” Report, University of Bonn. [0051]

Bayer, Christian, Ralph Luetticke, Lien Pham-Dao, and Volker Tjaden (2019), “Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk.” *Econometrica*, 87 (1), 255–290. [0053]

Berczik, Peter, Ruth Männer, Guillermo Marcus, Robi Banerje, Andreas Kugel, Ralf Klessen, and Gerhard Lienhart (2009), “Accelerating astrophysical particle simulations with programmable hardware (FPGA and GPU).” *Computer Science Research and Development*, 23, 231–239. [0050]

Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J. Sargent (2017), “Fiscal policy and debt management with incomplete markets.” *Quarterly Journal of Economics*, 132 (2), 617–663. [0051]

Biggs, Benjamin, Ian McInerney, Eric C. Kerrigan, and George A. Constantinides (2022), “High-level synthesis using the Julia language.” Technical report, Imperial College, London. [0069]

Bilal, Adrien (2021), “Solving heterogeneous agent models with the master equation.” Technical report, University of Chicago. [0051]

Brumm, Johannes and Simon Scheidegger (2017), “Using adaptive sparse grids to solve high-dimensional dynamic models.” *Econometrica*, 85 (5), 1575–1612. [0051]

Cai, Yongyang and Thomas S. Lontzek (2019), “The social cost of carbon with economic and climate risks.” *Journal of Political Economy*, 127 (6), 2684–2734. [0053]

Cheela, Bhagath, André DeHon, Jesús Fernández-Villaverde, and Alessandro Peri (2023), *A Beginner’s Guide to Programming FPGAs for Economics: An Introduction to Electrical Engineering Economics*. University of Pennsylvania. [0052, 0054, 0058, 0070]

Cheela, Bhagath, André DeHon, Jesús Fernández-Villaverde, and Alessandro Peri (2024), “Supplement to ‘Programming FPGAs for economics: An introduction to electrical engineering economics.’” *Quantitative Economics Supplemental Material*, 15, <https://doi.org/10.3982/QE2344>. [0054]

Childers, David (2018), “Solution of rational expectations models with function valued states.” Manuscript, Carnegie Mellon. [0051]

Corporation Intel (2021), *Intel® High Level Synthesis Compiler Pro Edition User Guide (UG20037)*. Intel Corporation. [0059]

Cruz Álvarez, José-Luis and Esteban Rossi-Hansberg (2021), “The economic geography of global warming.” Working Paper 28466, National Bureau of Economic Research. [0053]

Den Haan, Wouter J., Kenneth L. Judd, and Michel Juillard (2010), “Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty.” *Journal of Economic Dynamics and Control*, 34 (1), 1–3. [0051, 0053, 0054, 0055]

Den Haan, Wouter J. and Pontus Rendahl (2010), “Solving the incomplete markets model with aggregate uncertainty using explicit aggregation.” *Journal of Economic Dynamics and Control*, 34 (1), 69–78. [0051]

Duarte, Victor, Diogo Duarte, Julia Fonseca, and Alexis Montecinos (2019), “Benchmarking machine-learning software and hardware for quantitative economics.” *Journal of Economic Dynamics and Control*, 111, 103796. [0050]

Fernández-Villaverde, Jesús, Samuel Hurtado, and Galo Nuño (2019), “Financial frictions and the wealth distribution.” Working Paper 26302, National Bureau of Economic Research. [0051, 0082]

Fernández-Villaverde, Jesús and David Z. Valencia (2018), “A practical guide to parallelization in economics.” Working Paper 24561, National Bureau of Economic Research. [0050]

Frigo, Jan, Maya Gokhale, and Dominique Lavenier (2001), “Evaluation of the Streams-C C-to-FPGA compiler: An applications perspective.” In *Proceedings of the 2001 ACM/SIGDA Ninth International Symposium on Field Programmable Gate Arrays*, 134–140. [0051]

Goldberg, David (1991), “What every computer scientist should know about floating-point arithmetic.” *ACM Computing Surveys (CSUR)*, 23 (1), 5–48. [0068]

Herbordt, Martin C., Josh Model, Yongfeng Gu, Bharat Sukhwani, and Tom VanCourt (2006), “Single pass, blast-like, approximate string matching on FPGAs.” In *2006 14th Annual IEEE Symposium on Field-Programmable Custom Computing Machines*, 217–226. [0050]

Hoang, Dzung (1993), “Searching genetic databases on Splash 2.” In *[1993] Proceedings IEEE Workshop on FPGAs for Custom Computing Machines*, 185–191. [0050]

Hrica, James (2012), “Floating-point design with Vivado HLS.” Xilinx Application Note. [0068]

Judd, Kenneth L., Lillia Maliar, Serguei Maliar, and Inna Tsener (2017), “How to solve dynamic stochastic models computing expectations just once.” *Quantitative Economics*, 8 (3), 851–893. [0051]

Kadric, Edin, Paul Gurniak, and André DeHon (2016), “Accurate parallel floating-point accumulation.” *IEEE Transactions on Computers*, 65 (11), 3224–3238. [0068]

Kahou, Mahdi E., Jesús Fernández-Villaverde, Jesse Perla, and Arnav Sood (2021), “Exploiting symmetry in high-dimensional dynamic programming.” Working Paper 28981, National Bureau of Economic Research. [0051, 0082]

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante (2018), “Monetary policy according to HANK.” *American Economic Review*, 108 (3), 697–743. [0053]

Kapre, Nachiket and Andre DeHon (2009), “Performance comparison of single-precision SPICE model-evaluation on FPGA, GPU, cell, and multi-core processors.” In *International Conference on Field Programmable Logic and Applications (FPL)*, 65–72. [0069]

Krusell, Per and Anthony A. Smith (1998), “Income and wealth heterogeneity in the macroeconomy.” *Journal of Political Economy*, 106 (5), 867–896. [0049, 0051, 0052, 0053, 0054, 0055, 0058, 0064, 0066, 0068, 0069, 0070, 0073]

Krusell, Per and Anthony A. Smith (2022), “Climate change around the world.” Working Paper 30338, National Bureau of Economic Research. [0053]

Lo, Charles and Paul Chow (2016), “Model-based optimization of high level synthesis directives.” In *2016 26th International Conference on Field Programmable Logic and Applications (FPL)*, 1–10. [0069]

Maliar, Lillia, Serguei Maliar, and Fernando Valli (2010), “Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm.” *Journal of Economic Dynamics and Control*, 34 (1), 42–49. [0051, 0052, 0055, 0057, 0070, 0072, 0073]

Mertens, Thomas M. and Kenneth L. Judd (2018), “Solving an incomplete markets model with a large cross-section of agents.” *Journal of Economic Dynamics and Control*, 91, 349–368. [0051]

Nagurney, Anna (1996), “Parallel computation.” In *Handbook of Computational Economics 1*, 335–404. [0050]

Nurvitadhi, Eriko, Ganesh Venkatesh, Jaewoong Sim, Debbie Marr, Randy Huang, Jason Ong Gee Hock, Yeong Tat Liew, Krishnan Srivatsan, Duncan Moss, Suchit Subhaschandra, and Guy Boudoukh (2017), “Can FPGAs beat GPUs in accelerating next-generation deep neural networks?” In *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, 5–14. [0082]

Peri, Alessandro (2020), “A hardware approach to value function iteration.” *Journal of Economic Dynamics and Control*, 114, 1–18. [0050, 0051, 0058, 0059, 0074, 0077]

Pröhl, Elisabeth (2015), “Approximating equilibria with ex-post heterogeneity and aggregate risk.” Research Paper 17-63, Swiss Finance Institute. [0051]

Quenon, Alexandre and Victor R. G. da Silva (2021), “Towards higher-level synthesis and co-design with Python.” In *Proceedings of the Workshop on Languages, Tools, and Techniques for Accelerator Design (LATTE'21)*, ACM, New York, NY, USA. [0069]

Reiter, Michael (2009), “Solving heterogeneous-agent models by projection and perturbation.” *Journal of Economic Dynamics and Control*, 33 (3), 649–665. [0051]

Reiter, Michael (2010), “Solving the incomplete markets model with aggregate uncertainty by backward induction.” *Journal of Economic Dynamics and Control*, 34 (1), 28–35. [0051]

Snider, Greg (2002), “Performance-constrained pipelining of software loops onto reconfigurable hardware.” In *Proceedings of the 2002 ACM/SIGDA Tenth International Symposium on Field-Programmable Gate Arrays*, 177–186. [0051]

Whaley, R. Clint and Jack J. Dongarra (1998), “Automatically tuned linear algebra software.” In *Proceedings ACM International Conference on Supercomputing*, 1–27, Washington, DC, USA. [0069]

Winberry, Thomas (2018), “A method for solving and estimating heterogeneous agent macro models.” *Quantitative Economics*, 9 (3), 1123–1151. [0051]

Xilinx, Inc. (2020), *UG1145: Xilinx Vitis Unified Software Platform User Guide*. Xilinx, Inc. [0052, 0059]

Xilinx, Inc. (2021), *Overview of Arbitrary Precision Fixed-Point Data Types*. Xilinx. Accessed on 2023/11/02. [0061]

Yates, Randy (2009), “Fixed-point arithmetic: An introduction.” *Digital Signal Labs*, 81 (83), 198. [0061]

Young, Eric R. (2010), “Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations.” *Journal of Economic Dynamics and Control*, 34 (1), 36–41. [0051, 0081]

Young-Schultz, Tanner, Lothar Lilge, Stephen Brown, and Vaughn Betz (2020), “Using OpenCL to enable software-like development of an FPGA-accelerated biophotonic cancer treatment simulator.” In *Proceedings of the 2020 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, 86–96. [0050]

Zhao, Jieru, Liang Feng, Sharad Sinha, Wei Zhang, Yun Liang, and Bingsheng He (2020), “Performance modeling and directives optimization for high-level synthesis on FPGA.” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 39 (7), 1428–1441. [0069]

Co-editor Morten O. Ravn handled this manuscript.

Manuscript received 22 February, 2023; final version accepted 5 November, 2024; available online 5 November, 2024.

The replication package for this paper is available at <https://doi.org/10.5281/zenodo.14013936>. The Journal checked the data and codes included in the package for their ability to reproduce the results in the paper and approved online appendices. Given the highly demanding nature of the algorithms, the reproducibility checks were run on a simplified version of the code, which is also available in the replication package.