



Computational Challenges and Opportunities in Financial Services

Art Sedighi^(✉) and Doug Jacobson^(✉)

Department of Electrical and Computer Engineering, Iowa State University,
Ames, IA 50011, USA

{asedighi, dougj}@iastate.edu

Abstract. As one of the fastest growing areas of applied scientific computing, financial services uses high performance computing techniques to respond to both governmental regulatory bodies as well as to deal with a fast-paced business environment. Financial services industry is data driven and aims to resolve mathematical challenges to make sense out of data to solve complex problems in pricing, risk management, and portfolio optimization. These challenges are solved by financial institutions regularly, and the goal here is to provide a short survey of approaches and techniques used to solve these problems. Cloud is one of the areas of interest, since said challenges can benefit from the dynamicity and metered pricing of Cloud computing, plus being virtually limitless in scale. FPGA- and GPU-as-a-Service will also be explored as they are showing a great deal of benefit in solving such problems.

Keywords: Cloud · Financial services · Microsoft Azure · HPC · Machine learning · Risk calculation · Derivatives pricing · Portfolio optimization

1 Summary

Financial services is one of the fastest growing areas of applied scientific computing [1]. Computationally intensive problems solved by financial services require not only an extensive array of traditional x86 processors but have shown to benefit from other exotic and proprietary hardware such as GPU [2, 3], FPGA [4, 5] and even custom ASIC [6]. In addition to the heterogeneity of the infrastructure, different methods of computation have also been applied in financial services such machine learning, AI [7], various simulation methods [4, 5, 8–10], along with many different mathematical techniques from statistics, to numerical methods [11] to Stochastic Differential Equations [12]. This paper aims to provide a short survey of the top computationally intensive workloads in financial services, and the methods and techniques used to solve them. First, the top business questions that lead to computationally difficult problems will be outlined. This is followed by some methods used to solve these problems, albeit most use some form of heuristic method due to the difficulty entailed. We will then delve into the runtime aspect of these problems and what capabilities cloud, mainly Azure [13], provides vis-à-vis infrastructure capabilities to aide in solving said challenges.

2 Overview of Challenges in Financial Services

The majority of the complex problems in finance stem from one of the following three areas: (a) pricing, valuing and hedging of securities, (b) risk management [14] and (c) portfolio optimization [15].

Pricing of securities involved complex mathematical methods, most of which can only be solved using numerical methods or heuristically using methods like Monte Carlo simulations [3, 16–18]. The complexity arises from the fact that stock prices follow Brownian Motion (BM) in that they move in an irregular, non-differentiable, and non-smooth fashion [19]. To model stock prices, and as a result, the rest of the financial markets, the non-differentiability aspect of Brownian Motion requires the use of Stochastic Differential Equations (SDEs) [12]. SDEs discretize stock prices into infinitesimal pieces, model or simulate each component, and collect the result set.

The primary goal of risk management is to estimate the loss with a certain level of confidence (\$11.55–\$11.57 million with a confidence level of 99%, for example). The potential loss translates to capital requirement from a regulatory perspective, and to trading decision from a portfolio optimization perspective [14]. Risk management encompasses a number of different risk calculations, some of which are regulatory enforced [20–23]. The majority of risk calculations require vast computational power due to the complexity of the problem [9, 10, 24–29].

Portfolio optimization is affected by the preferences set by the investor, the constraints set by the investor, governmental bodies, regulations, taxes and many other decisions made to optimize the portfolio based on the aforementioned preferences and constraints [1]. Taxes and tax implications have particularly shown to have a great implication on portfolio optimization [30].

The aforementioned complexities are caused by mathematical challenges. Solving these problems using the set of tools and techniques at our disposal has received a great deal of attention over the years, and that research is still ongoing. Some of the details behind the mathematical problems will further be explored in the next section.

The ability to efficiently use the available technologies such as cloud, multi-core processors, FPGA's, GPUs, and purpose-build hardware in order to adequately and cost effectively solve complex mathematical methods will be covered in the later section of this paper.

3 Computational Challenges in Pricing and Valuing Securities

Derivative securities have payoff functions that are dependent on other securities, known as the underlying [31]. The primary use of derivatives is to hedge own risk or transfer the risk to others [14]. Pricing and valuing European options, American options, Asian options, Mortgage Backed Securities (MBS), credit derivatives, credit default swaps, and collateral debt obligations [11, 16, 32–35] are examples of derivative securities which present a computational challenge in finance. They all

depend on the value, price and performance of one or more underlying securities such as bonds or equities.

Monte Carlo (MC) method is used to simulate and determine options pricing. MC heavily depends of random numbers or pseudo-random numbers for pseudo-Monte Carlo [4] simulations. Generating truly random numbers at scale represents a computational challenge in financial services [36, 37]. The Computational complexity of the standard MC is $O(n^3)$ [38], with n representing the random number generated for the simulation. Reducing the computational complexity of MC to anything lower represents another challenge in financial services [39].

Pricing of exotic options [40] represent another computational challenge in finance [39]. Exotic options can be used in trading of different asset classes like commodities: coffee, corn, pork belly, and crude oil, along with equities, bonds and foreign exchange. All the aforementioned components need to be discretized individually. As different asset classes are added to the mix, discretizing will become a challenge, and prone to error [41]. As a result, the combinatorial of different choices of asset classes is explored, which itself introduces another computational challenge in financial services.

4 Computational Challenges in Risk Management

The core of risk management is to measure risk in the context of [42]:

- Market Risk as it relates to changes in prices of assets
- Value-at-Risk (VaR) as it relates to determining with a level of confidence the maximum loss over a period of time
- Credit Risk as it relates to the level of loss vis-à-vis probability of default of a payment of a loan or borrowing
- Operational Risk as it relates to the loss due to operational processes or an event, internal or externally caused
- Liquidity Risk as it relates to the inability of a financial institution to transform from one asset class to another

Monte Carlo (MC) or pseudo-Monte Carlo methods are generally used to calculate Value-at-Risk [9], and more generally speaking, to measure the risk posture in financial services. As mentioned in the prior section, the primary goal of risk management is to estimate the loss with a certain level of confidence. Calculating this loss function represents a challenge in financial services [39].

As is the case with heuristic approaches to solve complex problems, there is no single correct answer, so the simulations run many times with different parameters. Multi-level MC simulations are likely as the number of unknowns increase, and those unknowns need to be estimated as well, before final loss function estimation. Finding the balance between what is known as inner- and outer-Monte Carlo in terms of the number of unknowns and calculations represents another challenge in financial services. At the end, the results are collated and analyzed. This is a challenging problem by itself, although not computationally challenging but would fall under data challenges in financial services, as the correct result is not known, and further analysis is required to find the most likely result.

Computational challenges translate to longer running jobs. Even though tremendous amount of effort has gone into parallelizing these runs, there is still much more to be done. Risk calculations have normally been done overnight. Another computational challenge is to reduce the runtime of calculating risk, and moving towards what is known as intraday risk simulation and calculations [39, 43, 44].

5 Computational Challenges in Portfolio Optimization

Portfolio optimization involves a constraint-based allocation problem with an objective function that is based on the preferences set forth by the investor. As the number of investment options increase, so does the difficulty in the decision-making process. Computational finance has been used in portfolio optimization tasks [15], along with more advanced mathematical and optimization techniques [17]. Taxes and tax implication of portfolio optimization has added a level of difficulty to portfolio optimization as well [30]. The computational challenge with portfolio optimization is caused by the dynamicity aspect of the positions in a portfolio, and the quick response requirement that is elevated due to a vast pool of alternatives. Dynamic programming is used to model portfolio optimization [1, 15].

6 Data Challenges

Although the focus has been on computational challenges in finance, challenges in data are tied to the why and the how of the challenges presented.

Financial services has always been a data driven industry, but the conventional methods via which financial data has been managed have become a challenge for financial services. The telemetry generated by a financial transaction: seller information, buyer information, price, P&L, positions, along with industry-related telemetry such as interest rate, commodity prices, equity prices, bond rate, foreign exchange rate, have added to the level of interest coming from financial services. About 70% of US equity trades are machine-driven using High Frequency Trading (HFT) techniques. Due to advancement in technology, there is an exponential increase in data generated by recent HFT. This has translated to Ultrahigh Frequency Data (UHFD), and storing, accessing and making sense out of this data in a reasonable timeframe is a data challenge facing financial services [45]. The interest is shifting to UHFD as UHFD-based volatility models [46] have shown to have improved statistical efficiency, and be very useful in evaluation of daily volatility forecasts [47, 48]. With UHFD, financial services are now able to achieve accurate intraday volatility measures using machine learning techniques and statistical models [45, 49, 50].

However, computational challenges are now facing additional challenges brought upon by the increase in data volume.

7 Solutioning Options

Table 1 summarizes the different computational challenges outlined in this paper. These challenges are solved by financial institutions regularly, and the goal here is provide a survey of approaches to these challenges. This paper considers Microsoft Azure [13] as the cloud provider for this study. Azure offers a broad range of capabilities considered for this research including a range of Virtual machines [51], GPU's and FPGA's [52].

Table 1. Summary of challenges and solutioning options

Challenges	Bottleneck	Cloud	GPU	FPGA
Random number generation at scale	Computation and distribution	Yes [65]	Yes [2]	Yes [4, 39, 66]
Reducing $O(n^3)$ complexity of Monte Carlo	Computation and model creation	Yes [63]	Yes [63]	N/A
Discretization of [pricing] data that follows Brownian Motion	Model creation	Not needed or No Data Available	Yes [3]	Yes [39]
Multi-level Monte-Carlo (inner-loop vs. outer loop optimization)	Computation	Not needed or No Data Available	Yes [67, 68]	Yes [39, 67, 68]
Reducing VaR runtime to enable intraday risk calculation	Computation	Yes [69]	Yes [68]	Yes [39, 68]
Dynamic portfolio optimization that is tax efficient	Computation	Yes [1, 15, 17]	No Data Available	No Data Available
Portfolio optimization in highly dynamic portfolio positions	Computation and data distribution	Yes [1, 15, 17]	No Data Available	No Data Available
Ultra-high frequency data challenge	Data lifecycle management	Yes [60]	N/A	N/A

Properties of cloud [53–55] have shown to benefit the large-scale needs of financial services. General Purpose GPU (GPGPU) computation can aide in solving problems with high parallelization needs. Field Programmable Gate Arrays (FPGA) based computing has the capability to be customized to a specific problem, and as such to dramatically increase performance.

The overall goal for any option presented is to allow for intra-day risk calculation to occur. Intra-day risk calculation is hindered by strict time requirement. If calculation needs to occur to measure risk across the entire portfolio before the decision on the next transaction to take place, then the timing allotted to do such calculation is bound by the speed of business. In other words, there is no limit on the need to speed up these calculations. This requirement sets the basis for much of the large-scale deployments of compute in financial services.

7.1 Cloud

The properties of cloud such as metered usage, and dynamic provisioning is useful in most cases, specially more so in cases where quick response is needed that would require a dynamic buildout of a new environment [56]. The dynamic provisioning of resources can reduce computation time as it allows virtually unlimited scalability. The dynamicity of cloud to rapidly scale up to 1000's of VM's [57, 58] can reduce the need to spent time creating complex models and allow for rapid testing. This can potentially reduce the overall runtime of the calculations and allow for intra-day risk models [59] to run. Fast VMs coupled with the availability of high-speed networking and technology like Infiniband [57, 58] can reduce the interprocess communication delays, which in turn can assist with the task of random number distribution. To tackle the storage bottleneck introduced by UHFD on the cloud, a cloud-based high-speed storage environment would be desirable [60].

Machine learning and optimization techniques are often used for portfolio optimization [15, 61, 62]. A new approach that can potentially replace traditional Monte Carlo simulations with Neural Networks has shown to reduce calculation times exponentially [63].

In additions to the traditional VM-based machine learning deployments, machine learning and optimization models are now finding their way on both Cloud-based FPGAs [62], and GPUs [63, 64].

7.2 GPU

GPU-based infrastructure has been used to speed up risk and pricing calculation in financial services [2, 3, 70]. The many-core capability of GPU's have the ability to run potentially 1000's of tasks in parallel, thus reducing the overall runtime. Some cloud-based GPU's are capable of ~ 3 Teraflops of double-precision operation across ~ 5000 cores [71].

As mentioned before, one of the main challenges of financial services is options pricing, and more specifically, exotic options pricing. GPU's have been used to be able to do both very efficiently [72–74]. GPUs have been used develop a Neural Network-based approach to replace Monte Carlo simulations in pricing derivatives [63]. GPUs have the capability to provide raw computational power, with little to change to the method by which these models are solved. FPGA's, however, have the capability to introduce a new paradigm in solving computational challenges in financial services.

7.3 FPGA

FPGA's allow for an alternate method of solving computationally challenging problems. Random number generation is inherent in FPGA's and thus easier to compute as demonstrated in [66]. Quasi-random numbers have also been tested and proven to function at great speeds on FPGA's [4]. Putting an entire Monte Carlo simulation on an FPGA has also been demonstrated in [5, 75] with a good degree of success.

Azure has demonstrated the performance gain achievable with their FPGA-based SmartNIC [76] network cards. The authors demonstrate a <15 ms VM-VM TCP latencies and 40+ Gbps throughput. The reconfigurability of the FPGA-based network cards proved to be a valuable feature as changes and bugs could simply be deployed in real-time. FPGA's [and GPU's] are used as CPU offloads, in that tasks are off-loaded to the FPGA in order to reduce the load on the CPU cores. This does require code change, as the process needs to be FPGA aware to utilize its capabilities. Even though the development costs are higher, cloud providers are still providing FPGA computing -as-a-Service as part of their service portfolio.

8 Conclusion

Financial services need to respond to both governmental regulatory bodies as well as deal with fast paced business environments. The combination of the two has made the financial services sector a hotbed for challenging problems and innovation. Pricing, portfolio optimization, and risk management are the main three areas from which challenges are derived. Computation and data challenges outlined in this paper have received a great deal of attention from the research community. Computational challenges can benefit from the dynamicity and metered pricing of Cloud computing, plus being virtually limitless in scale. FPGA- and GPU-as-a-Service are also becoming available as they are showing a great deal of benefit in solving such problems. The next step for this research is to evaluate specific benefits that stems from using Cloud, FPGA and GPU in solving these challenges.

References

1. Haugh, M.B., Lo, A.W.: Computational challenges in portfolio management. *Comput. Sci. Eng.* 3(3), 54 (2001)
2. Grauer-Gray, S., et al.: Accelerating financial applications on the GPU. In: *Proceedings of the 6th Workshop on General Purpose Processor Using Graphics Processing Units*. ACM (2013)
3. Solomon, S., Thulasiram, R.K., Thulasiraman, P.: Option pricing on the GPU. In: *2010 IEEE 12th International Conference on High Performance Computing and Communications (HPCC)*. IEEE (2010)
4. Banks, S., Beadling, P., Ferencz, A.: FPGA implementation of pseudo random number generators for Monte Carlo methods in quantitative finance. In: *2008 International Conference on Reconfigurable Computing and FPGAs*. IEEE (2008)

5. Woods, N.A., VanCourt, T.: FPGA acceleration of quasi-Monte Carlo in finance. In: 2008 International Conference on Field Programmable Logic and Applications. IEEE (2008)
6. Pottathuparambil, R., et al.: Low-latency FPGA based financial data feed handler. In: 2011 IEEE 19th Annual International Symposium on Field-Programmable Custom Computing Machines. IEEE (2011)
7. Krollner, B., Vanstone, B.J., Finnie, G.R.: Financial time series forecasting with machine learning techniques: a survey. In: ESANN (2010)
8. Eckhardt, R.: Stan Ulam, John von Neumann, and the Monte Carlo method. *Los Alamos Sci.* **15**(131–136), 30 (1987)
9. Glasserman, P., Heidelberger, P., Shahabuddin, P.: Efficient Monte Carlo methods for value-at-risk (2010)
10. Tezuka, S., et al.: Monte Carlo grid for financial risk management. *Future Gener. Comput. Syst.* **21**(5), 811–821 (2005)
11. Gobet, E.: Advanced Monte Carlo methods for barrier and related exotic options. In: *Handbook of Numerical Analysis*, pp. 497–528. Elsevier (2009)
12. Kloeden, P.E., Platen, E.: *Numerical Solution of Stochastic Differential Equations*, vol. 23. Springer, Berlin (2013). <https://doi.org/10.1007/978-3-662-12616-5>
13. Microsoft Azure (2018). <http://azure.microsoft.com>. Accessed 12 Feb 2018
14. Staum, J.: Monte Carlo computation in finance. In: L'Ecuyer, P., Owen, A. (eds.) *Monte Carlo and Quasi-Monte Carlo Methods*. Springer, Berlin (2009)
15. Joseph, T.: Computational financing techniques and fundamental challenges in portfolio optimization. *IOSR J. Hum. Soc. Sci.* **9**(6), 51–58 (2013)
16. Black, F., Scholes, M.: The pricing of options and corporate liabilities. *J. Polit. Econ.* **81**(3), 637–654 (1973)
17. Korn, R., Korn, E.: *Option Pricing and Portfolio Optimization: Modern Methods of Financial Mathematics*, vol. 31. American Mathematical Society (2001)
18. Korn, R., Müller, S.: Binomial Trees in Option Pricing—History, Practical Applications and Recent Developments. In: Devroye, L., Karasözen, B., Kohler, M., Korn, R. (eds.) *Recent Developments in Applied Probability and Statistics*, pp. 59–77. Springer, Berlin (2010). https://doi.org/10.1007/978-3-7908-2598-5_3
19. Karatzas, I., Shreve, S.E.: *Brownian Motion and Stochastic Calculus*. Springer, New York (1991)
20. Committee, B.: *Basel III: a global regulatory framework for more resilient banks and banking systems*. Basel Committee on Banking Supervision, Basel (2010)
21. Gleeson, S.: *International Regulation of Banking: Basel II: Capital and Risk Requirements*. OUP Catalogue (2010)
22. Hakenes, H., Schnabel, I.: Bank size and risk-taking under Basel II. *J. Bank. Finance* **35**(6), 1436–1449 (2011)
23. Tarullo, D.K.: *Banking on Basel: The Future of International Financial Regulation*. Peterson Institute (2008)
24. Alexander, C.: Volatility and correlation: measurement, models and applications. *Risk Manag. Anal.* **1**, 125–171 (1998)
25. Brummelhuis, R., et al.: Principal component value at risk. *Math. Finance* **12**(1), 23–43 (2002)
26. Chong, J., Keutzer, K., Dixon, M.F.: *Acceleration of Market Value-at-Risk Estimation*. Available at SSRN 1576402 (2009)
27. Giot, P.: Market risk models for intraday data. *Eur. J. Finance* **11**(4), 309–324 (2005)
28. Jorion, P.: *Value at Risk*. McGraw-Hill, New York (1997)
29. McNeil, A.J., Frey, R., Embrechts, P.: *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press (2015)

30. Bertsimas, D., Lo, A.W.: Optimal control of execution costs. *J. Financ. Mark.* **1**(1), 1–50 (1998)
31. Bodie, Z., et al.: *Investments*. McGraw-Hill Education (2015)
32. Heston, S.L.: A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Rev. Financ. Stud.* **6**(2), 327–343 (1993)
33. Giesecke, K.: An overview of credit derivatives. Available at SSRN 1307880 (2009)
34. Giesecke, K.: Portfolio credit risk: top-down versus bottom-up approaches. *Front. Quant. Finance*, 251 (2009)
35. Fabozzi, F.J.: *The Handbook of Mortgage-Backed Securities*. Oxford University Press (2016)
36. Gentle, J.E.: *Random Number Generation and Monte Carlo Methods*. Springer, Berlin (2006). <https://doi.org/10.1007/b97336>
37. Niederreiter, H.: Quasi-Monte Carlo methods and pseudo-random numbers. *Bull. Am. Math. Soc.* **84**(6), 957–1041 (1978)
38. Anderson, D.F., Higham, D.J., Sun, Y.: Computational complexity analysis for Monte Carlo approximations of classically scaled population processes. *Multiscale Model. Simul.* **16**(3), 1206–1226 (2018)
39. Desmettre, S., Korn, R.: 10 computational challenges in finance. In: De Schryver, C. (ed.) *FPGA Based Accelerators for Financial Applications*, pp. 1–31. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-15407-7_1
40. Zhang, P.G.: *Exotic options: a guide to second generation options*. World Scientific (1998)
41. Dempster, M., Hutton, J.: Fast numerical valuation of American, exotic and complex options. *Appl. Math. Finance* **4**(1), 1–20 (1997)
42. Pan, S.-Q.: A survey of financial risk measurement. In: 6th International Conference on Management Science and Management Innovation (MSMI 2019). Atlantis Press (2019)
43. Sedighi, A., Deng, Y., Zhang, P.: Fairness of task scheduling in high performance computing environments. *Scalable Comput.: Pract. Exp.* **15**(3), 273–285 (2014)
44. Sedighi, A., Smith, M.: *Fair Scheduling in High Performance Computing Environments*. Springer, Cham (2019). <https://doi.org/10.1007/978-3-030-14568-2>
45. Seth, T., Chaudhary, V.: *Big Data in Finance* (2015)
46. Barndorff-Nielsen, O.E., Shephard, N.: Power and bipower variation with stochastic volatility and jumps. *J. Financ. Econom.* **2**(1), 1–37 (2004)
47. Bollerslev, T., Wright, J.H.: High-frequency data, frequency domain inference, and volatility forecasting. *Rev. Econ. Stat.* **83**(4), 596–602 (2001)
48. Grammig, J., Wellner, M.: Modeling the interdependence of volatility and inter-transaction duration processes. *J. Econom.* **106**(2), 369–400 (2002)
49. Comte, F., Renault, E.: Long memory in continuous-time stochastic volatility models. *Math. Finance* **8**(4), 291–323 (1998)
50. McAleer, M., Medeiros, M.C.: Realized volatility: a review. *Econom. Rev.* **27**(1–3), 10–45 (2008)
51. High performance compute VM sizes. Virtual Machine Documentation (2019). <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sizes-hpc>. Accessed 22 July 2019
52. What are field-programmable gate arrays (FPGA) (2019). <https://docs.microsoft.com/en-us/azure/machine-learning/service/concept-accelerate-with-fpgas>. Accessed 22 July 2019
53. Armbrust, M., et al.: *Above the Clouds: A Berkeley View of Cloud Computing* (2009)
54. Avram, M.-G.: Advantages and challenges of adopting cloud computing from an enterprise perspective. *Procedia Technol.* **12**, 529–534 (2014)
55. Armbrust, M., et al.: A view of cloud computing. *Commun. ACM* **53**(4), 50–58 (2010)
56. Smith, D.M.: *Cloud computing primer for 2016*. Gartner Inc., Stamford (2016)

57. Azure HC-series Virtual Machines cross 20,000 cores for HPC workloads (2019). <https://azure.microsoft.com/en-us/blog/azure-hc-series-virtual-machines-crosses-20000-cores-for-hpc-workloads/>. Accessed 22 July 2019
58. Working with large virtual machine scale sets (2019). <https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-placement-groups>. Accessed 22 July 2019
59. Enabling the financial services risk lifecycle with Azure and R. (2019). <https://docs.microsoft.com/en-us/azure/industry/financial/fsi-risk-modeling>. Accessed 22 July 2019
60. Cray in Azure (2019). <https://azure.microsoft.com/en-us/solutions/high-performance-computing/cray/>. Accessed 22 July 2019
61. What is axiomaBlue? (2019). <https://www.axioma.com/products/axiomablue/>. Accessed 22 July 2019
62. Deploy a model as a web service on an FPGA with Azure Machine Learning service (2019). <https://docs.microsoft.com/en-us/azure/machine-learning/service/how-to-deploy-fpga-web-service>. Accessed 22 July 2019
63. Ferguson, R., Green, A.D.: Deeply learning derivatives. Available at SSRN 3244821 (2018)
64. Deploy a deep learning model for inference with GPU (2019). <https://docs.microsoft.com/en-us/azure/machine-learning/service/how-to-deploy-inferencing-gpus>. Accessed 22 July 2019
65. Kerrigan, B., Chen, Y.: A study of entropy sources in cloud computers: random number generation on cloud hosts. In: Kotenko, I., Skormin, V. (eds.) MMM-ACNS 2012. LNCS, vol. 7531, pp. 286–298. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33704-8_24
66. Yap, A.Y.: Information Systems for Global Financial Markets: Emerging Developments and Effects: Emerging Developments and Effects. IGI Global (2011)
67. Tian, X., Benkrid, K.: High-performance quasi-monte carlo financial simulation: FPGA vs. GPP vs. GPU. *ACM Trans. Reconfigurable Technol. Syst. (TRETs)* **3**(4), 26 (2010)
68. Singla, N., et al.: Financial Monte Carlo simulation on architecturally diverse systems. In: 2008 Workshop on High Performance Computational Finance. IEEE (2008)
69. Kim, H., et al.: Online risk analytics on the cloud. In: Proceedings of the 2009 9th IEEE/ACM International Symposium on Cluster Computing and the Grid. IEEE Computer Society (2009)
70. Qiu, M., et al.: Data transfer minimization for financial derivative pricing using Monte Carlo simulation with GPU in 5G. *Int. J. Commun Syst* **29**(16), 2364–2374 (2016)
71. Azure N-Series VMs and NVIDIA GPUs in the Cloud (2016). <https://buildazure.com/azure-n-series-vm-and-nvidia-gpus-in-the-cloud/>. Accessed 31 July 2019
72. Bernemann, A., Schreyer, R., Spanderen, K.: Accelerating exotic option pricing and model calibration using GPUs. Available at SSRN 1753596 (2011)
73. Gaikwad, A., Toke, I.M.: GPU based sparse grid technique for solving multidimensional options pricing PDEs. In: Proceedings of the 2nd Workshop on High Performance Computational Finance. ACM (2009)
74. Abbas-Turki, L.A., Lapeyre, B.: American options pricing on multi-core graphic cards. In: 2009 International Conference on Business Intelligence and Financial Engineering. IEEE (2009)
75. De Schryver, C. (ed.): FPGA Based Accelerators for Financial Applications. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-15407-7>
76. Firestone, D., et al.: Azure accelerated networking: SmartNICs in the public cloud. In: 15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2018) (2018)