Journal of Computational Physics **179,** 698–703 (2002)
doi:10.1006/jcph.2002.7083

Moore's Law and Numerical Modeling

V. R. Voller and F. Porté-Agel

Civil Engineering, University of Minnesota, Minneapolis, Minnesota 55455
E-mail: volle001@umn.edu

Received August 23, 2001; revised April 8, 2002

An estimate of the rate of increase in numerical simulation grid sizes with time is obtained by counting the grids (measured in terms of number of node points) reported in the nine volumes of an established proceedings on the numerical modeling of solidification phenomena dating back to 1980. It is shown that the largest grids used in a given year increase at a rate consistent with the well-known Moore's law on computing power, i.e., the number of nodes in the grids double every 18 months. From this observation, approximate bounds on the available grid size in a current year are established. This approximation is used to provide projections as to when, assuming Moore's law continues to hold, direct simulations of physical phenomena, which resolve to the smallest scale present, will be achievable.

INTRODUCTION

In the mid 1960s Gordon Moore, the co-founder of Intel, made the observation that computer power, P, measured by the number of transistors that could be fit onto a chip, doubled once every 1.5 years [1]. This law, which has performed extremely well over the proceeding 30 or so years, can be stated in mathematical terms as

$$P = A2^{0.6667Y}, \tag{1}$$

where A is the computer power at the reference year $Y = 0$.

Major benefactors of an increase in computer power are those researchers working on the development of numerical models of engineering and physical phenomena. Increased power allows a refinement of the spatial and temporal approximations of the governing equations and also allows researchers to accurately account for the wide range of length and time scales present.

The objects of this article are to

1. Demonstrate that the increase in the size of numerical models, measured in terms of the number of node points used in the spatial discretization, matches Moore's law [1].

0021-9991/02 $35.00

2. Discuss how this finding can be used to make predictions about the future of numerical modeling, in particular, predictions of when direct numerical simulations of physical phenomena will be possible.

DATA SOURCES

The main data source, for tracking the increase in grid sizes, is the nine volumes from the conference "Modeling of Casting Welding and Solidification Processes" [2–10], which cover the time period of 1980 ($Y = 0$) to 2000 ($Y = 20$). These proceedings have a number of key attributes:

1. The topic area is well focused.
2. The topic requires sophisticated computational models of heat and mass transfer processes and phenomena.
3. The conference is well established.
4. The conference is relatively selective ($\sim$140 papers per volume).
5. The central core of research groups involved has remained reasonably consistent over the years.
6. The proceedings is a flagship in the area, where the current state of the art in modeling of casting and welding is well represented.

To provide additional points for comparison two one-off conferences—held in 1987 and 1999 respectively—in the closely related area of materials process modeling [11, 12] are also analyzed.

METHOD AND RESULTS

The measure of the numerical grid used is in terms of the total number of grid points, and throughout this work a term such as "large grid size" refers to a computation that uses a large number of grid points. In selecting the grid sizes from the papers in the proceedings, the following rules were used:

1. Modeling results for the given grid size had to be reported—in many cases large grid sizes were projected but results were not presented to support that calculations were possible.
2. Only the largest grid size from each paper was recorded.
3. The grid size can be reported as number of "nodes" or number of "elements." In a typical application, the factor between nodes and elements is between 0.5 and 2 and as such no attempt has been made to differentiate between nodes and elements.
4. In some cases, a figure of the grid used was supplied but no explicit numbers were provided. If the grid was two dimensional and the figure easy to read, then the grid points were counted by hand.
5. Some models employed adaptive meshing in which the mesh size evolved with the calculation; counts on these were ignored.

Although this exercise is a little bit like Mendel's assistants counting beans it should provide a consistent picture of how grid sizes have increased through the years in response to increased computer power.

TABLE I
Grid Sizes in Node Points

Year	Papers reporting grids	Grid size (page number)	Grid size (page number)	Grid size (page number)
1980 [2]	7	900 (52)	2,160 (204)	3,454 (143)
1983 [3]	8	861 (6)	954 (426)	32,400 (450)
1986 [4]	10	2,500 (176)	2,510 (502)	4,900 (217)
1987 [11]	11	1,759 (1130)	5,400 (435)	6,875 (418)
1988 [5]	20	46,875 (590)	57,621 (848)	152,361 (789)
1991 [6]	19	180,000 (28)	308,000 (731)	1,068,144 (774)
1993 [7]	19	738,000 (395)	1,000,000 (26)	2,193,750 (395)
1995 [8]	24	72,464 (217)	400,000 (360)	413,952 (750)
1998 [9]	38	1,728,000 (579)	6,349,500 (767)	34,409,821 (440)
1999 [12]	20	470,000 (98)	2,097,152 (67)	50,000,000 (357)
2000 [10]	32	5,000,000 (321)	10,396,350 (369)	120,000,000 (234)

The three largest recorded grid sizes (in ascending order) are tabulated in Table I. This table gives the year, number of papers reporting a grid size, and also matches the page number in the proceeding volume to the recorded grid size.

Figure 1 plots the logs of the three largest grid sizes from each volume against year. As points of reference two realizations of Moore's law, Eq. (1), are plotted using values of $A = 100$ and 10,000. Although, as might be expected, there is a fairly high degree of scatter, Fig. 1 clearly shows that increases in grid sizes do match Moore's law. For a given year, reasonable bounds on the maximum grid size used, are given by

$$100\ \ 2^{0.6667(\text{Year}-1980)} < N(t) < 10000\ 2^{0.6667(\text{Year}-1980)} \tag{2}$$

Further, an overall best fit to the data in Table I results in the single equation

$$N(t) = 691\ 2^{0.697(\text{Year}-1980)}. \tag{3}$$

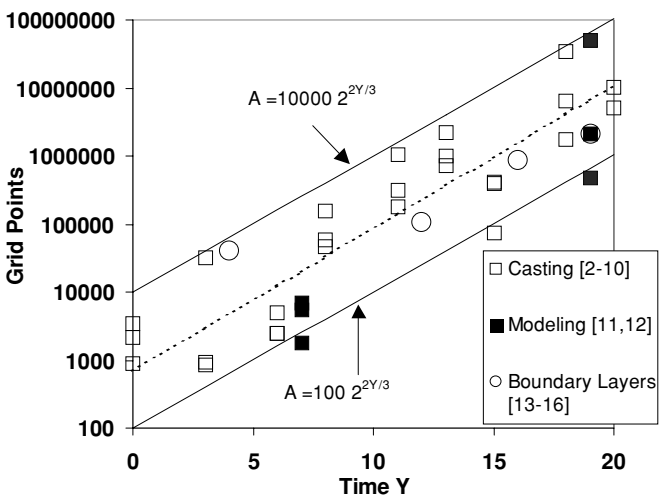

FIG. 1. Log of three largest grid sizes from each volume plotted against year.

This equation is shown as a dashed line in Fig. 1. Note the exponent, $p = 0.697$, is very close to the value of 0.667 required by Moore's law.

Although the bounds in Eq. (2) are specific to the area of materials modeling, they do appear to be appropriate for other numerical modeling fields. The open circles in Fig. 1 correspond to a sampling of grid sizes used in large eddy simulations of atmospheric boundary layers [13–16]. In this case, the search through all available grids is not as rigorous as for the casting case. Nevertheless, the sampling does indicate that grids used in the modeling of atmospheric boundary layers do fall inside the bounds given by Eq. (2).

DISCUSSION

The essence of modeling casting and welding processes is to arrive at, via the solution of heat and fluid flow equations, a complete description of the chemical and structural composition in a cast product. A direct simulation of such a process will require a domain size at the process level ($\sim$1 m) that will have to resolve phenomena down to the length scale of the solid–liquid interface. Typically, the smallest length scale will be the order of a dendrite tip radius, which could approach 1 μm [17]. In some cases, however, simulations may need to resolve to the lattice spacing ($\sim$1 nm) [17]. In the former case, 10^6 nodes will be needed in each dimension, in the later 10^9. The direct simulation of a 1 m^3 casting resolving phenomena to the lattice length scale will need 10^{27} nodes.

A similar spread of length scales is needed in the direct numerical simulation of atmospheric boundary layers [18]. In this case a typical domain size will be $\sim$10 km but calculations will have to be resolved to the smallest turbulent length scale, that is, on the order of 1 mm; such calculations will need 10^7 node points in each direction.

Assuming generality and continued validity of the bound in Eq. (2), projections can be made as to when various direct simulations of physical phenomena will be possible; see Table II. From this table it can be observed that complete three-dimensional direct numerical simulations of casting and turbulence phenomena cannot be expected to occur until well into this century and, in some cases, not until the turn of the century.

TABLE II
Expected Year ($\pm$5) That the Given Direct Simulation Will Be Possible If Grid Size Increases Are Bound by Eq. (2)

Simulation	Domain length scale	Resolution length scale	Grid points required	Expected year ($\pm$5 years)
2-D casting	0.1 m	1 μm (dendrite tip)	10^{10}	2015
2-D casting	1 m	1 μm (dendrite tip)	10^{12}	2025
3-D casting	0.1 m	1 μm (dendrite tip)	10^{15}	2040
Boundary layer	100 m	1 mm	10^{15}	2040
2-D casting	0.1 m	1 nm (lattice space)	10^{16}	2045
3-D casting	1 m	1 μm (dendrite tip)	10^{18}	2055
2-D casting	1 m	1 nm (lattice space)	10^{18}	2055
Boundary layer	1 km	1 mm	10^{18}	2055
Boundary layer	10 km	1 mm	10^{21}	2070
3-D casting	0.1 m	1 nm (lattice space)	10^{24}	2085
3-D casting	1 m	1 nm (lattice space)	10^{27}	2100

In the case of modeling solidification, this situation can be improved by the use of adaptive meshing in the vicinity of the solid–liquid interface [19, 20] or the coupling of conventional numerical solutions with Monte Carlo [21] and cellular automata [22]. In using adaptive grids in solidification, modeling computational requirements have been shown to scale with the arc-length of the solid–liquid interface [19]. In one recent example of adaptive meshing, the number of nodes in a 3-D calculation was reduced from 4×10^6 to 1.2×10^5 nodes [20]; this effectively reduced the dimension for the problem from 3 to 2.3. If this grid saving scales to larger problems it will reduce the projected time for a complete 3-D casting simulation by $\sim$30 years (2100 to 2070).

CONCLUDING REMARKS

This article has demonstrated that numerical grid sizes used in the simulation of physical phenomena follow Moore's law by doubling in size every 18 months. It remains an open question as to why such a scaling should occur. The observed scaling could suggest that the willingness to wait for results has remained constant over time. Such an interpretation, however, is confused if one considers issues related to calculation complexity and the impact of parallel architectures. In this light, the authors cannot see a clear scientific methodology that would uncover specific reasons why the scaling should hold. We leave it to the reader to speculate.

The study of the specific research area of solidification modeling has led to an explicit approximation (Eq. (2)) on how grid sizes (measured in terms of the number of nodes) would increase in time if Moore's law continued to hold. This model suggests that it may not be until the later part of this century that computing will be powerful enough to tackle direct numerical simulations of phenomena such as metal solidification and atmospheric boundary layers. For the foreseeable future, full-scale modeling of physical processes will have to be based on grids with spacing significantly larger than the smallest length scales of the underlying phenomena. Hence, research efforts aimed at

1. using volume averaging and sub grid models to capture the average effects of small scale phenomena in larger scale process models, and/or
2. the continued development of adaptive grid strategies will continue to be a key area of research in computational physics.

ACKNOWLEDGMENTS

This work was supported in part by the Minnesota Supercomputing Institute. FP-A was supported by the National Science Foundation through Grant EAR-0094200 and by NASA through Grant NAG5-10569. The authors appreciate the useful comments from Jon Dantzig, University of Illinois, and Christoph Beckermann, University of Iowa, on earlier drafts. The authors also thank the reviewers for insightful comments. In particular, the reviewer who suggested the inclusion of a single fit line for the data in Table I and the discussion on possible reasons why the Moore's Law scaling should hold.

REFERENCES

1. Gordon E. Moore, Cramming more components onto integrated circuits, *Electronica* **38** (1965).
2. H. D. Brody and D. Apelian, *Modeling of Casting and Welding and Processes* (TMS, Warrendale, 1981).
3. J. A. Dantzig and J. T. Berry, *Modeling of Casting and Welding Processes—II* (TMS, Warrendale, 1984).

4. S. Kou and R. Mehrabian, *Modeling of Casting and Welding Processes—III* (TMS, Warrendale, 1986).
5. A. F. Giamei and G. J. Abbaschian, *Modeling of Casting and Welding Processes—IV* (TMS, Warrendale, 1988).
6. M. Rappaz, M. R. Ozgu, and K. W. Mahin, *Modeling of Casting, Welding and Advanced Solidification Processes—V* (TMS, Warrendale, 1991).
7. T. S. Piwonka, V. R. Voller, and L. Katgerman, *Modeling of Casting, Welding and Advanced Solidification Processes—VI* (TMS, Warrendale, 1993).
8. M. Cross and J. Campbell, *Modeling of Casting, Welding and Advanced Solidification Processes—VII* (TMS, Warrendale, 1995).
9. B. G. Thomas and C. Beckermann, *Modeling of Casting, Welding and Advanced Solidification Processes—VIII* (TMS, Warrendale, 1998).
10. P. R. Sahm, P. N. Hansen, and J. G. Conley, *Modeling of Casting, Welding and Advanced Solidification Processes—IX* (Shaker Verlag, Aachen, 2000).
11. J. Szekely, L. B. Hales, H. Henein, N. Jarrett, K. Rajamani, and I. Samarsekera, *Mathematical Modeling of Materials Processing Operations* (TMS, Warrendale, 1987).
12. N. El-Kaddah, D. G. C. Robertson, S. T. Johansen, and V. R. Voller, *Fluid Flow Phenomena in Metals Processing* (TMS, Warrendale, 1999).
13. C.-H. Moeng, A large-eddy simulation model for the study of planetary boundary-layer turbulence, *J. Atmos. Sci.* **41**, 2052 (1984).
14. F. T. M. Nieuwstadt, P. J. Mason, C.-H. Moeng, and U. Schumann, Large-eddy simulation of the convective boundary layer: A comparison of four computer codes, in *Turbulent Shear Flows 8.* edited by F. Durst *et al.* (Springer-Verlag, Berlin, 1992).
15. Ayotte, K. W., P. P. Sullivan, A. Andren, S. C. Doney, A. A. M. Holtslag, W. G. Large, J. C. McWilliams, C.-H. Moeng, M. J. Otte, J. J. Tribbia, and J. C. Wyngaard, An evaluation of neutral and convective planetary boundary layer parameterizations relative to large eddy simulation, *Bound. Layer Meteor.* **79**, 131 (1996).
16. A. Juneja and J. G. Brasseur, Characteristics of subgrid-resolved-scale dynamics in anisotropic turbulence, with application to rough-wall boundary layers, *Phys. Fluids* **11**, 3054 (1999).
17. J. A. Warren and W. J. Boettinger, Predictions of dendritic growth and microsegregation patterns in a binary alloy using the phase-field method, *Acta Metall. Mater* **43**, 689 (1995).
18. R. B. Stull, *An Introduction to Boundary Layer Meteorology* (Kluwer Academic, Dordrecht, Norwell, MA, 1988).
19. N. Provatas, N. Goldenfeld, and J. Dantzig, Adaptive mesh refinement computation of solidification microstructures using dynamic data structures, *J. Comp Phys.* **148**, 265 (1999).
20. J. Dantzig, Private communication, June 27, 2001.
21. M. Plapp and A. Karma, Multiscale random-walk algorithm for simulating interfacial pattern formation, *Phys. Rev. Lett.* **84**, 1740 (2000).
22. Ch.-A. Gandin, J.-L. Desbiolles, M. Rappaz, and Ph. Thévoz, A Three-dimensional cellular automaton—Finite element model for the prediction of solidification grain structures, *Met. Mater. Trans.* **30A**, 3153 (1999).