

COPING WITH UNCERTAINTY AND STRESS: A PARALLEL COMPUTATION APPROACH*

Selim G. Akl[†]

May 8, 2003

Abstract

This paper is concerned with computations whose characteristics are akin to certain unique phenomena that occur in different domains of science. We are particularly interested in systems whose parameters are altered unpredictably whenever one of these parameters is measured or modified. Examples of such computational environments include those in which *Heisenberg's uncertainty principle* of quantum physics is witnessed, or those in which *Le Châtelier's principle* of chemical systems under stress manifests itself. A study of these systems uncovers computations that are inherently parallel *in the strong sense*, meaning that they are efficiently executed in parallel, but impossible to carry out sequentially.

Keywords: equilibrium, parallel computation, real time, uncertainty.

1 Introduction

The purpose of this paper is to bring to light the fact that in some special computational environments, a computation can succeed if and only if it is

*This research was supported by the Natural Sciences and Engineering Research Council of Canada.

[†]School of Computing, Queen's University, Kingston, Ontario, Canada K7L 3N6, Email: akl@cs.queensu.ca, Tel: (613) 533-6062, Fax: (613) 533-6513.

performed in parallel. In these environments, it is the laws of nature that prevail rather than human-imposed computational circumstances or conditions on the computation. Specifically, it is the principles governing such fields as biology, chemistry, and physics, that cause the inevitable failure of any sequential approach to solving the problem at hand, while at the same time allowing a parallel approach to succeed. A typical example of such principles is the uncertainty involved in measuring several related parameters of a physical system. Another principle expresses the way in which the components of a system in equilibrium react when subjected to outside stress.

The characteristics of the computational environment considered in this paper are summarized as follows. Consider a physical system \mathcal{S} . Thus, \mathcal{S} may be a system studied by biologists (for example, an ecosystem), or one maintained by engineers (for example, a nuclear reactor). It is required to perform the following operations on \mathcal{S} . First, a set of n parameters are to be measured. The values of these parameters are then used to compute new values for a second set of n parameters of the same system \mathcal{S} . Finally, these computed values are applied to \mathcal{S} . One property of the physical system is that measuring or setting a parameter modifies the values of any number of other parameters in the set unpredictably. This modification is forced by the scientific principles that are germane to the computational environment. Neither humans nor computers have any control over these natural phenomena. We show in this paper that, under these conditions, a parallel approach succeeds in performing the task while a sequential approach fails.

It is important to contrast the situation described in the previous paragraph with former analyses of the improvement (in speed or quality of the solution) obtained when performing a computation in parallel. Those analyses include ones performed within conventional paradigms (such as, for example, when all the data required by a computation are available at the outset), as well as analyses performed within unconventional paradigms (such as, for example, when the data arrive in real time and the results must be delivered by a certain strict deadline). They also include those analyses yielding improvements (in speed or quality) that are sublinear, linear, or even superlinear in the number of processors used on the parallel computer. Unlike the case with \mathcal{S} , where nature dictates the behavior of the system, the conditions governing the computational environment in traditional analyses are, in a fashion, fully decided by the human operator and the model of computation used. For example, in a real-time computation, if it is deemed that the arrival rate of the data is too high, it is possible for the people in charge of the computation

to slow down the arrival rate, or to extend the deadline by which the solution is to be delivered, or to use a faster computer, and so on. Not so with \mathcal{S} .

The remainder of this paper is organized as follows. The characteristics of a physical system \mathcal{S} , as well as the computational problem to be solved on it, are described in Section 2. Sections 3 and 4, respectively, present the models of computation to be used for designing different approaches to the problem, and the assumptions we make in order to analyze these approaches. Parallel and sequential solutions are derived and analyzed in Section 5. The discussion in Section 6 is intended to provide examples of phenomena from various disciplines in the natural and physical sciences; each of these phenomena is a specialization of a property of \mathcal{S} . Some thoughts as to the consequences of our results are provided in Section 7.

Throughout the paper a *time unit* is understood to be the time required by a computer to perform a basic computational step like reading or writing or performing a fundamental arithmetic or logical operation (for example, addition, comparison, and so on) on a constant number of data of fixed size. It is important to keep in mind that the length of a time unit is not an absolute quantity. Instead, the duration of a time unit is defined in terms of the speed of the available processor.

2 Computational Problem

A physical system \mathcal{S} possesses the following characteristics:

1. For $n > 1$, the system possesses two sets, each of n parameters (or properties), namely, q_1, q_2, \dots, q_n and r_1, r_2, \dots, r_n , respectively. Each of these parameters is a physical quantity (such as, for example, temperature, humidity, density, pressure, electric charge, and so on). These quantities can be measured and/or controlled independently, each at a given discrete location (or point) within \mathcal{S} . Henceforth, q_i , $1 \leq i \leq n$, is used to denote a parameter as well as the discrete location at which this parameter is measured and/or controlled. The same holds for r_i , $1 \leq i \leq n$. It should be noted that the following discussion also applies in the special case where the two sets of parameters q_1, q_2, \dots, q_n and r_1, r_2, \dots, r_n are identical. It is also worth pointing out that the two sets of parameters are assumed to be of the same size only for simplicity of exposition.

2. The system is in a state of equilibrium, meaning that the values q_1, q_2, \dots, q_n satisfy a certain global condition $\mathcal{C}(q_1, q_2, \dots, q_n)$.
3. At regular intervals, the state of the physical system is to be recorded. In other words, the values q_1, q_2, \dots, q_n are to be measured at a given moment in time where $\mathcal{C}(q_1, q_2, \dots, q_n)$ is satisfied. Each interval has a duration of \mathcal{T} time units; that is, the state of the system is measured every \mathcal{T} time units.
4. If the values q_1, q_2, \dots, q_n are measured *one by one*, each separately and independently of the others, this disturbs the equilibrium of the system. Specifically, suppose (without loss of generality) that the parameters q_1, q_2, \dots, q_{i-1} have already been measured, for some $i, 1 < i < n$. Now, when q_i is subsequently measured, *at least one* other value $q_j, 1 \leq j \leq n$ and $j \neq i$, will change unpredictably shortly thereafter (within one time unit), such that $\mathcal{C}(q_1, q_2, \dots, q_n)$ is no longer satisfied. Most importantly, the values of $q_{i+1}, q_{i+2}, \dots, q_n$, none of which has yet been registered, may be altered irreparably.
5. If the values q_1, q_2, \dots, q_n can be measured correctly, such that $\mathcal{C}(q_1, q_2, \dots, q_n)$ holds, then new values for r_1, r_2, \dots, r_n are to be computed and applied to \mathcal{S} in order to bring q_1, q_2, \dots, q_n to a new desired state of equilibrium.
6. This computation of the new values for r_1, r_2, \dots, r_n must be done as quickly as possible. This is because the time during which the system is in a state of disequilibrium is to be minimized (for safety reasons, for example), and furthermore the system is not to be allowed to reach a new undesirable state of equilibrium on its own (for integrity reasons, for example). To be specific, a new (desired) state of equilibrium must be reached within \mathcal{T} time units (in other words, before the state of the system is next measured).
7. Suppose that new values have been computed for r_1, r_2, \dots, r_n . Setting these parameters to their new values *consecutively*, causes all parameter values to be arbitrarily modified each time a parameter is set. For example, suppose (without loss of generality) that the parameters r_1, r_2, \dots, r_{i-1} , have received their new values, for some $i, 1 < i < n$. Now, setting r_i to its new value independently, may as a consequence

cause any or all of the r_j , $1 \leq j \leq n$ and $j \neq i$, to be altered in an unpredictable way, shortly thereafter (within one time unit). Most importantly, the values of r_1, r_2, \dots, r_{i-1} , which have already been set, may be altered irreparably.

8. Once the new values of the parameters r_1, r_2, \dots, r_n have been applied successfully, the parameters q_1, q_2, \dots, q_n acquire new values satisfying $\mathcal{C}(q_1, q_2, \dots, q_n)$.

The foregoing requirements, namely, that the parameters of \mathcal{S} be measured, new parameters be computed and applied, and a new (desired) state of equilibrium be reached, all within \mathcal{T} time units, suggest that the computation needs to be performed in real time [4, 8, 13]. A fundamental property of *real-time computation* is that certain operations must be performed by specified *deadlines*. Thus, one or more of the following conditions may be imposed:

1. Each received input (or set of inputs) must be processed within a certain time after its arrival.
2. Each output (or set of outputs) must be returned within a certain time after the arrival of the corresponding input (or set of inputs).

Often, as in this paper, all deadlines are *tight*, that is, they are measured in terms of one time interval, and they are *firm*, meaning that missing a deadline causes the computation to fail [2].

3 Models of Computation

Two models of computation are presented for addressing the problem of Section 2, one sequential, the other parallel.

3.1 Sequential model

The first model consists of one processor along with some memory and internal registers. This is the standard model used in conventional algorithm analysis. Despite its familiarity, however, there are many instances of this model. For definiteness, we assume in what follows that the Random Access Machine (RAM) version of the model [1] is used.

3.2 Parallel model

The second model consists of several processors, each executing its algorithm. These processors work simultaneously on the solution to a computational problem. Here, again, there are many options from which to choose. For convenience, we adopt the Parallel Random Access Machine (PRAM), a model that is quite well-known [1]. Here, n processors share a common memory through which they communicate. The processors work synchronously and execute the same algorithm.

Both models described use the same type of processors. In particular, the RAM processor has the same computational capabilities as each PRAM processor. Furthermore, both the RAM processor and each PRAM processor run at the same speed (which we assume to be the maximum speed possible theoretically).

4 Simplifying Assumptions

In order to perform a concrete analysis of the different solutions to the computational problem outlined in Section 2, we assume in what follows that each of the following operations requires one time unit:

1. Measuring one parameter q_i , $1 \leq i \leq n$
2. Computing one parameter r_i , $1 \leq i \leq n$
3. Applying one parameter r_i , $1 \leq i \leq n$, to \mathcal{S} .

Furthermore, once the new values of the parameters r_1, r_2, \dots, r_n have been applied to \mathcal{S} , the system requires one additional time unit to reach a new state of equilibrium. It follows that the shortest \mathcal{T} can be is four time units; we therefore assume that $\mathcal{T} = 4$.

5 Solutions

Two approaches are now described for addressing the problem of Section 2, namely, to measure the state of \mathcal{S} while in equilibrium, thus disturbing the latter, then restoring it.

5.1 Sequential Approach

The RAM will measure *one* of the values (q_1 , for example) and by so doing disturb the equilibrium, thus losing all hope of recording the state of the system within the given time interval. Any value read afterwards will not satisfy $\mathcal{C}(q_1, q_2, \dots, q_n)$.

Similarly, the sequential approach cannot update the parameters of \mathcal{S} properly: Once r_1 has received its new value, setting r_2 disturbs r_1 unpredictably.

For the sake of argument, suppose that the sequential approach can do the following:

1. Measure all of q_1, q_2, \dots, q_n simultaneously (with, for example, n sensors connected to a single processor, all of them measuring at the same time and filling up a memory array of length n with the values q_1, q_2, \dots, q_n).
2. Update all of r_1, r_2, \dots, r_n simultaneously (with, for example, n controllers connected to the single processor, all of them activated at the same time once the new values have been computed).

Despite these assumptions, the sequential approach would still be incapable of computing the new parameters within \mathcal{T} time units, as it obviously requires n time units to do so. This is true even when n takes its smallest value of 2, since in this case 5 time units are needed to reach a new equilibrium, and $5 > \mathcal{T}$.

5.2 Parallel Approach

The PRAM, by contrast, will measure *all* the parameters q_1, q_2, \dots, q_n simultaneously (one value per processor), and therefore obtain an accurate reading of the state of the system within the given time frame. Consequently,

1. A snapshot of the state of the system that satisfies $\mathcal{C}(q_1, q_2, \dots, q_n)$ has been obtained.
2. The new parameters r_1, r_2, \dots, r_n can be computed in parallel in one time unit (one value per processor).
3. These new parameters can also be applied to the system simultaneously (one value per processor) in one time unit.

Following the resetting of the parameters r_1, r_2, \dots, r_n , the parameters q_1, q_2, \dots, q_n in turn settle at their new values thus reaching a new equilibrium. The entire process concludes within \mathcal{T} time units successfully.

6 Discussion

A computational problem was described in connection with a physical system \mathcal{S} . To recap, \mathcal{S} is initially in a state of equilibrium. Some of its parameters must be measured; however, measurement disturbs the equilibrium. The system is not to be left out of equilibrium for too long. A new set of values for (the same or) other parameters of \mathcal{S} are computed and these parameters are set to their new values. The system enters a new desired state of equilibrium.

As shown in Section 5, the process above can be carried out successfully in parallel, but not sequentially. This result is due to two properties possessed by \mathcal{S} :

1. Measuring one parameter of \mathcal{S} affects the other parameters unpredictably.
2. Modifying one parameter of \mathcal{S} affects the other parameters unpredictably.

These two properties are reminiscent of a number of well-known principles that manifest themselves in such diverse fields as the physical and natural sciences, engineering, and sociology, to cite but a few. Some of these are listed below, grouped for our purposes into two classes.

6.1 Uncertainty in measurement

Phenomena in this class occur in systems where measuring one parameter of the system affects, interferes with, or even precludes the subsequent measurement of another parameter of the system. It is important to emphasize that the kind of uncertainty of concern here is in no way due to any errors that may be introduced by an imprecise or not sufficiently accurate measuring apparatus.

1. In quantum mechanics, *Heisenberg's uncertainty principle* puts a limit on our ability to measure simultaneously pairs of 'complementary' variables. Thus, the *position* and *momentum* of a subatomic particle, or

the *energy* of a particle in a certain state and the *time* during which that state existed, cannot be defined at the same time to arbitrary accuracy [5]. In fact, what this principle says is that once *one* of the two variables is measured (however accurately, but independently of the other), the act of measuring itself introduces a disturbance that affects the value of the *other* variable. For example, suppose that at a given moment in time t_0 the position p_0 of an electron is measured. Assume further that it is also desired to determine the electron's momentum m_0 at time t_0 . When the momentum is measured, however, the value obtained is not m_0 , as it would have been changed by the previous act of measuring p_0 . On a macroscopic scale, quantum phenomena occur in systems exhibiting properties such as superconductivity, superfluidity, and magnetism [14]. An example of a large-scale system of this kind is the superconducting quantum interference device (SQUID) proposed as a possible building block for quantum computers [16].

2. In digital signal processing the *uncertainty principle* is exhibited when conducting a Fourier analysis. Complete resolution of a signal is possible either in the time domain t or the frequency domain w , but not both simultaneously. This is due to the fact that the Fourier transform is computed using e^{iwt} : Since the product wt must remain constant, narrowing a function in one domain, causes it to be wider in the other [7, 12]. For example, a pure sinusoidal wave has no time resolution, as it possesses nonzero components over the infinitely long time axis. Its Fourier transform, on the other hand, has excellent frequency resolution: It is an impulse function with a single positive frequency component. By contrast, an impulse (or *delta*) function has only one value in the time domain, and hence excellent resolution. Its Fourier transform is the constant function with nonzero values for all frequencies and hence no resolution.

Other examples in this class include image processing, sampling theory, spectrum estimation, image coding, and filter design [17].

6.2 Reaction to stress

Phenomena in this class arise in systems where modifying the value of a parameter causes a change in the value of another parameter. In response

to stress from the outside, the system automatically reacts so as to relieve the stress. Newton's third law of motion ("For every action there is an equal and opposite reaction") is a good way to characterize these phenomena.

1. In chemistry, *Le Châtelier's principle* states that if a system at equilibrium is subjected to a stress, the system will shift to a new equilibrium in an attempt to reduce the stress. The term *stress* depends on the system under consideration. Typically, stress means a change in pressure, temperature, or concentration [10]. For example, consider a container holding gases in equilibrium. Decreasing (increasing) the volume of the container leads to the pressure inside the container increasing (decreasing); in response to this external stress the system favors the process that produces the least (most) molecules of gas. Similarly, when the temperature is increased (decreased), the system responds by favoring the process that uses up (produces) heat energy. Finally, if the concentration of a component on the left (right) side of the equilibrium is decreased (increased), the system's automatic response is to favor the reaction that increases (decreases) the concentration of components on the left (right) side.
2. In biology, the *homeostatic principle* is concerned with the behavior displayed by an organism to which stress has been applied [11, 15]. An automatic mechanism known as *homeostasis* counteracts external influences in order to maintain the equilibrium necessary for survival, at all levels of organization in living systems. Thus, at the molecular level, homeostasis regulates the amount of enzymes required in metabolism. At the cellular level, it controls the rate of division in cell populations. Finally, at the organismic level, it helps maintain steady levels of temperature, water, nutrients, energy, and oxygen. Examples of homeostatic mechanisms are the sensations of hunger and thirst. In humans, sweating and flushing are automatic responses to heating, while shivering and reducing blood circulation to the skin are automatic responses to chilling. Homeostasis is also seen as playing a role in maintaining population levels (animals and their prey), as well as steady state conditions in the Earth's environment.

Systems with similar behavior are also found in cybernetics, economics, and the social sciences [9].

Each of the phenomena discussed in Sections 6.1 and 6.2 typically involves *two* variables in equilibrium. Measuring or setting one of the variables has an impact on the value of the other variable. The system \mathcal{S} , however, involves *several* variables (two or more). In that sense, its properties, as listed at the beginning of this section, are extensions of these phenomena.

7 Conclusion

Can a parallel computer with n processors solve a computational problem more than n times *faster* than a sequential computer? Can it solve it more than n times *better*? In the latter case, the definition of quality depends on the problem domain. Thus, in combinatorial optimization it measures how close an approximate solution is to the optimal one, in numerical analysis it expresses the accuracy of a computation, and so on. In [3], several computational paradigms are described that offer affirmative answers to the above questions. Concrete examples are presented in which the improvement in speed or quality of the answer is *superlinear* in the number of processors used by the parallel computer (meaning that, for example, the improvement is on the order of n^x , where x is a constant greater than 1). Furthermore, the improvement is consistent and provable. All examples are characterized by the presence of one or several input streams arriving in real time. Here we need to underline the fact that each of these examples, in essence, represents an existence proof that superlinear behavior can indeed be achieved in parallel computation.

In this paper, an attempt is made to go one step further. A hypothetical system \mathcal{S} is conceived. The system has n measurable parameters q_i and n controllable parameters r_i . Furthermore, \mathcal{S} is in a state of equilibrium. Measuring one of the q_i disturbs the equilibrium and, as a consequence, causes all measurable parameters (but most importantly those *not yet* measured) to be altered unpredictably. Similarly, setting one of the r_i causes all controllable parameters (but most importantly those *already* set) to be altered unpredictably. It is shown that the task of measuring the q_i , computing new values for the r_i , and setting the latter before the system settles into an undesirable state of equilibrium can be performed in parallel but not se-

quentially. We have therefore established, in theory, that there indeed exist problems that are *inherently parallel* in the strong sense. These problems are not merely solvable efficiently in parallel—their solution can *only* be obtained in parallel.

Evidence was provided that the properties of \mathcal{S} occur in many areas of scientific study, including living organisms, albeit often in specialized and reduced forms. One field currently receiving a fair deal of attention is *complexity* (not to be confused with computational complexity). It endeavors to study complex systems, that is, systems that are self-regulating, or adaptive, or whose behavior may be described as nonlinear, and so on [6]. It is here perhaps that additional systems possessing the properties of \mathcal{S} are found. The discovery of such complex systems in the real world, and the study of their properties, will have profound consequences on the science of computing, both in theory and in practice.

References

- [1] S.G. Akl, *Parallel Computation: Models and Methods*, Prentice Hall, Upper Saddle River, New Jersey, 1997.
- [2] S.G. Akl, Parallel real-time computation: Sometimes quantity means quality, *Proceedings of the International Symposium on Parallel Architectures, Algorithms and Networks*, Dallas, Texas, December 2000, pp. 2 - 11.
- [3] S.G. Akl, Superlinear performance in real-time parallel computation, *Proceedings of the Thirteenth Conference on Parallel and Distributed Computing and Systems*, Anaheim, California, August 2001, pp. 505 - 514.
- [4] A. Bestavros and V. Fay-Wolfe, Eds., *Real-Time Database and Information Systems*, Kluwer Academic Publishers, Boston, 1997.
- [5] B.H. Bransden and C.J. Joachain, *Quantum Mechanics*, Pearson Education, Harlow (Essex), England, 2000.
- [6] J.L. Casti, *Complexification*, HarperCollins, New York, 1998.
- [7] D. Gabor, Theory of communication, *Proceedings of the Institute of Electrical Engineers*, Vol. 93, No. 26, 1946, 420–441.

- [8] H.W. Lawson, *Parallel Processing in Industrial Real-Time Applications*, Prentice Hall, Englewood Cliffs, New Jersey, 1992.
- [9] R. Lewin, *Complexity*, The University of Chicago Press, Chicago, 1999.
- [10] L.W. Potts, *Quantitative Analysis*, Harper & Row, New York, 1987.
- [11] D.B. Pribor, *Functional Homeostasis: Basis for Understanding Stress*, Kendall Hunt, Dubuque, Iowa, 1986.
- [12] C.E. Shannon, Communication in the presence of noise, *Proceedings of the IRE*, Vol. 37, 1949, 10–21.
- [13] M. Thorin, *Real-Time Transaction Processing*, Macmillan, London, 1992.
- [14] D.R. Tilley and J. Tilley, *Superfluidity and Superconductivity*, John Wiley and Sons, New York, 1974.
- [15] P.A. Trojan, *Ecosystem Homeostasis*, Dr. W. Junk, Dordrecht, The Netherlands, 1984.
- [16] T. Van Duzer and C.W. Turner, *Principles of Superconductive Devices and Circuits*, Prentice Hall, Upper Saddle River, New Jersey, 1999.
- [17] R. Wilson and G.H. Granlund, The uncertainty principle in image processing, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-6, No. 6, 1984, 758–767.