ties perpendicular to the magnetic field. The frequency of the lower hybrid waves is between the gyrofrequencies of the electrons (ω_{ce}) and the ions (ω_{ci}) which means that these waves can be in simultaneous Cherenkov resonance with the relatively slow but unmagnetized ions perpendicular to the magnetic field and fast magnetized (hence magnetic field aligned) electrons. Cherenkov resonance occurs when the phase velocity of the wave and the particle velocity are equal; under these conditions strong interaction between the waves and particles is possible and results in energy transfer from the wave to the particle or vice versa. The lower hybrid waves provide the intermediary step in transferring energy between the ions and electrons.
5. M. J. Mumma *et al.*, *Science* **272**, 1310 (1996).
6. D. Krankowsky *et al.*, *Nature* **321**, 326 (1986).
7. H. S. Hudson, W.-H. Ip, D. A. Mendis, *Planet. Space Sci.* **29**, 1373 (1981).
8. J. B. McBride, E. Ott, P. B. Jay, J. H. Orens, *Phys. Fluids* **157**, 2367 (1972). A two stream instability results when two charged particle populations traveling in opposite directions interact.
9. M. K. Wallis and R. S. B. Ong, *Planet. Space Sci.* **23**, 713 (1975). A more accurate calculation based on the analysis of the solar wind dynamics, mass-loaded by the picked-up cometary ions lead to the same formula for the ion density.
10. D. A. Mendis, H. L. F. Houpis, M. L. Marconi, *Physics of Comets Fundamentals of Cosmic Physics* (1985), vol. 10.
11. L. D. Landau, *J. Phys. USSR* **10**, 25 (1946); F. F. Chen, *Introduction to Plasma Physics and Controlled Fusion* (Plenum, New York, 1984), vol. 1, p. 240.
12. V. D. Shapiro and V. I. Shevchenko, *Sov. Sci. Rev. E, Astrophys. Space Phys.* **6**, 425 (1988).
13. D. F. Post, R. V. Jensen, C. B. Tarter, W. H. Grasberger, W. A. Lokke, *Princeton Plasma Physics Laboratory Report PPPL-1352* (1977).
14. J. M. Dawson, in *Fusion*, E. Teller, Ed. (Academic Press, New York, 1981), p. 465.
15. J. W. Chamberlain, *Physics of the Aurora and Airglow* (Academic Press, New York, 1961).
16. This work was supported in part by NSF grant PH-9319198;003 and NASA NAGW-1502.

21 June 1996; accepted 17 October 1996

An Economics Approach to Hard Computational Problems

Bernardo A. Huberman, Rajan M. Lukose, Tad Hogg

A general method for combining existing algorithms into new programs that are unequivocally preferable to any of the component algorithms is presented. This method, based on notions of risk in economics, offers a computational portfolio design procedure that can be used for a wide range of problems. Tested by solving a canonical NP-complete problem, the method can be used for problems ranging from the combinatorics of DNA sequencing to the completion of tasks in environments with resource contention, such as the World Wide Web.

Extremely hard computational problems are pervasive in fields ranging from molecular biology to physics and operations research. Examples include determining the most probable arrangement of cloned fragments of a DNA sequence (*1*), the global minima of complicated energy functions in physical and chemical systems (*2*), and the shortest path visiting a given set of cities (*3*), to name a few. Because of the combinatorics involved, their solution times grow exponentially with the size of the problem (a basic trait of the so-called NP-complete problems), making it impossible to solve very large instances in reasonable times (*4*).

In response to this difficulty, a number of efficient heuristic algorithms have been developed. These algorithms, although not always guaranteed to produce a good solution or to finish in a reasonable time, often provide satisfactory answers fairly quickly. In practice, their performance varies greatly from one problem instance to another. In many cases, the heuristics involve randomized algorithms (*5*), giving rise to performance variability even across repeated trials on a single problem instance.

In addition to combinatorial search problems, there are many other computational situations where performance varies from one trial to another. For example, programs operating in large distributed systems or interacting with the physical world can have unpredictable performance because of changes in their environment. A familiar example is the action of retrieving a particular page on the World Wide Web. In this case, the usual network congestion leads to a variability in the time required to retrieve the page, raising the dilemma of whether to restart the process or wait.

In all of these cases, the unpredictable variation in performance can be characterized by a distribution describing the probability of obtaining each possible performance value. The mean or expected values of these distributions are usually used as an overall measure of quality (*6*–*9*). We point out, however, that expected performance is not the only relevant measure of the quality of an algorithm. The variance of a performance distribution also affects the quality of an algorithm because it determines how likely it is that a particular run's performance will deviate from the expected one. This variance implies that there is an inherent risk associated with the use of such an algorithm, a risk that, in analogy with the economic literature, we will identify with the standard deviation of its performance distribution (*10*).

Risk is an important additional characteristic of algorithms because one may be willing to settle for a lower average performance in exchange for increased certainty in obtaining a reasonable answer. This situation is often encountered in economics when trying to maximize a utility that has an associated risk. It is usually dealt with by constructing mixed strategies that have desired risk and performance (*11*). In analogy with this approach, we here present a widely applicable method for constructing "portfolios" that combine different programs in such a way that a whole range of performance and risk characteristics become available. Significantly, some of these portfolios are unequivocally preferable to any of the individual component algorithms running alone. We verify these results experimentally on graph-coloring, a canonical NP-complete problem, and by constructing a restart strategy for access to pages on the Web.

To illustrate this method, consider a simple portfolio of two Las Vegas algorithms, which, by definition, always produce a correct solution to a problem but with a distribution of solution times (*5*). Let t_1 and t_2 denote the random variables, which have distributions of solution times $p_1(t)$ and $p_2(t)$. For simplicity, we focus on the case of discrete distributions, although our method applies to continuous distributions as well. The portfolio is constructed simply by letting both algorithms run concurrently but independently on a serial computer. Let f_1 denote the fraction of clock cycles allocated to algorithm 1 and $f_2 = 1 - f_1$ be the fraction allocated to the other. As soon as one of the algorithms finds a solution, the run terminates. Thus, the solution time t is a random variable related to those of the individual algorithms by

$$t = \min\left(\frac{t_1}{f_1}, \frac{t_2}{f_2}\right) \tag{1}$$

The resulting portfolio algorithm is characterized by the probability distribution $p(t)$ that it finishes in a particular time t. This probability is given by the probability that both constituent algorithms finish in time $\geq t$ minus the probability that both algorithms finish in time $> t$

$$p(t) = \left[\sum_{t' \geq f_1 t} p_1(t')\right]\left[\sum_{t' \geq f_2 t} p_2(t')\right] - \left[\sum_{t' > f_1 t} p_1(t')\right]\left[\sum_{t' > f_2 t} p_2(t')\right] \tag{2}$$

Dynamics of Computation Group, Xerox Palo Alto Research Center, Palo Alto, CA 94304, USA.

Each value of f_1 therefore corresponds to a distribution whose mean and variance can be calculated. By varying f_1 from 0 to 1 and defining the risk σ as $\sqrt{\text{var}(t)}$, a curve in the plane of expected solution time $\langle t \rangle$ versus risk can be determined.

The simplest algorithm that exhibits interesting behavior is one with two possible solution times; that is, the algorithm can solve a problem in time $t = 1$ with probability p or in time $t = T$ with probability $1 - p$. By letting two independent instances of such an algorithm run, we obtain, using the above results, the risk versus expected time curve shown in Fig. 1 for particular values of p and T.

There are several features of this curve that are worth noting. First, endpoint A corresponds to the algorithm running alone ($f_1 = 1$ or 0), and point B corresponds to both algorithms sharing computer time equally ($f_1 = f_2 = ½$). Second, there exists a regime, the efficient frontier (indicated by the solid segment of the curve), defined by the fact that for every point on the curve, there is at least one point on the efficient frontier that is always preferable, that is, has a lower risk or higher performance, or both. Once this efficient frontier is determined, one can choose the desired performance-risk combination on it and calculate the corresponding fraction of computer cycles to be allocated to algorithm 1. This calculation can be done by plotting both the expected solution time and the risk as a function of f_1 (Fig. 2).

We point out that the efficient frontier for this distribution will persist as T is reduced from 10 to 3.5. The lower bound corresponds to a value of the ratio $\sigma/\langle t \rangle = 0.65$, known as the Sharpe ratio in the finance literature, which indicates the kinds of distributions that can be used in this approach.

Because Eq. 2 applies to any discrete distribution and can readily be generalized to continuous ones, this procedure also applies to more complicated situations. For example, by extending Eq. 2 to include the minimum of more than two random variables, the portfolio approach can be used in the case of many algorithms, each with their own fractional share of computer cycles. In that case, varying the fractions allocated to each algorithm produces a two-dimensional region in the risk-expected time plane, rather than a single curve. The efficient frontier is then a subset of the boundary of that region.

To test this portfolio approach with more realistic distribution functions, we experimented with the often studied NP-complete problem of graph-coloring (*12–14*). This problem consists of a graph (a collection of points or nodes, some of which are connected by straight edges), a specified number of colors, and the requirement to assign a color to each node in the graph such that no two nodes linked by an edge have the same color. The average degree of the graph γ (the average number of edges coming from a node in the graph) is a convenient parameter describing the amount of constraint in the problem and determines the typical behavior of a variety of search methods (*15*). Here, we focus on the case of three-coloring with $\gamma = 3.5$, which has been shown to exhibit a large variance of finishing times over the class of such graphs (*16*).

We used a complete, depth-first backtracking search based on the Brelaz heuristic (*13*), which assigns the most constrained nodes first (that is, those with the most distinctly colored neighbors), breaking ties by choosing nodes with the most uncolored neighbors (with any remaining ties broken randomly). For each node, the smallest numbered color consistent with the previous assignments was chosen first, with successive choices made when the search was forced to backtrack. This search method is guaranteed to terminate eventually having correctly found a possible coloring or having concluded that no such coloring exists for the graph. Thus, it is a Las Vegas algorithm.

We first produced a probability distribution of solution times (Fig. 3) by running the Brelaz heuristic algorithm repeatedly on a graph-coloring instance (*16*). Because the distribution shows a large variability in solution time, we expect that our method will make accessible points in the risk-expected solution time plane that are preferable to the one point accessible by running the heuristic by itself.

Application of Eq. 2 to two independent instances of this algorithm produces the expected solution time versus risk curve shown in Fig. 4. The functional similarity between this curve and that in Fig. 1 is apparent, confirming the predictions of our simplified example. In this case, the efficient frontier is in the range $0.013 < f_1 < 0.060$. The low end of this range locates the point of minimum risk, and the high end locates the point of maximum performance. The smallness of f_1 in this case shows that only a slight "mixing" of strategies is required to provide benefits because of the relatively high probability that the algorithm will find the solution fairly quickly. These results confirm the benefits to be accrued, for as the graph shows, performance can be increased by about 30% at reduced risk when a combined algorithm is used.

We also tested whether an efficient portfolio constructed for this particular graph

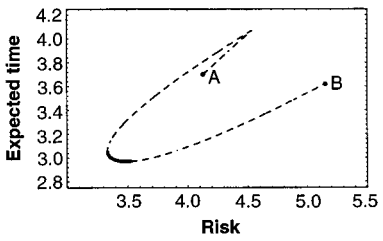

Fig. 1. Expected solution time versus risk for the case of two independent algorithms with identical discrete bimodal distributions. Here $p = 0.7$ and $T = 10$. Each point on the curve represents the risk and performance of the "portfolio" algorithm when the fraction f_1 ranges from 0 to 1: endpoint A marks $f_1 = 1 - f_2 = 0$ or 1, and endpoint B marks $f_1 = f_2 = ½$. The solid segment corresponds to the efficient frontier. Note that the shorter the expected solution time, the higher the performance, and vice versa.

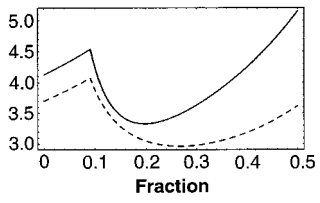

Fig. 2. Risk (solid) and expected solution time (dashed) versus fraction f_1 for the situation depicted in Fig. 1.

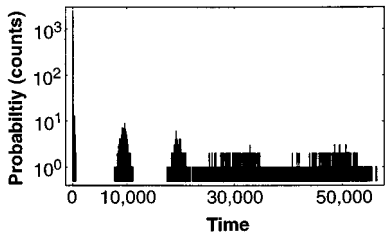

Fig. 3. A measured probability distribution for the Brelaz heuristic used on a particular 100-node graph-coloring problem with connectivity $\gamma = 3.5$. The distribution is the result of searching the graph 10,000 times with different random seeds.

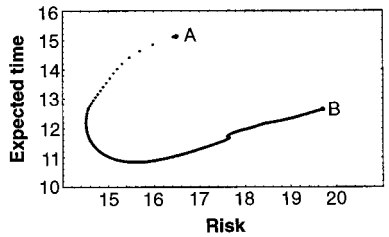

Fig. 4. Experimental performance versus risk curve for graph-coloring by a search method based on the Brelaz heuristic. This curve was produced by varying f_1 from 0 to 1. For each fraction, the risk and expected solution time was calculated from the distribution in Fig. 3 by using Eq. 2. Both axes are shown in units of 10^3 time steps. The expected performance and risk for the heuristic running alone is indicated by point A; point B indicates $f_1 = ½$.

would be effective on other graphs whose distributions of solution times were unknown. Using a portfolio with $f_1 = 0.013$ on 20 randomly chosen graphs yielded average speed and risk improvements of 22 and 10%, respectively.

Having established the utility of a portfolio approach when the component algorithms have highly variable performance, we point out that independent studies of a variety of NP-complete problems, using very different search algorithms, have discovered similar distributions in the context of phase transitions in search (*15*, *17*). It has also been pointed out that any algorithm that performs a depth-first backtracking search through a hierarchical tree will have a highly extended distribution of performance because early, high-level choices can decide immediately whether a particular run will take a short time or a much longer time (*7*, *9*). This variability in performance suggests that it is possible to predict when a particular instance of a heuristic is likely to have the right properties for this approach to be useful, thus making it very general in terms of applications.

So far we have assumed that the component algorithms are completely independent of each other and do not communicate. They can be thought of as "competing" with one another for machine resources. However, allowing for cooperation or dependencies among the individual algorithms while they are running simultaneously can improve performance (*18*, *19*). This possibility raises the interesting question of the extent to which our economics approach to portfolio algorithms can also benefit from cooperation. Basically, cooperation will introduce statistical correlations between the performance of the individual algorithms, and we will accordingly define the correlation between them as

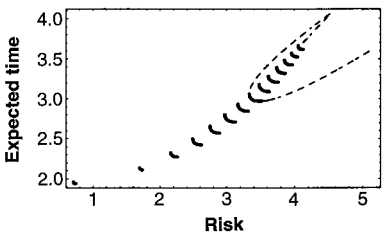

Fig. 5. Effect of cooperation among algorithms. For the case of two algorithms with the discrete bimodal distribution of solution times studied above, the correlation between the two distributions was varied to model the effect of cooperation between them. For values of ρ ranging from −0.42 at the lower left to 0.42 at the upper right, the efficient frontiers are plotted. The entire risk-expected solution time curve corresponding to ρ = 0 from Fig. 1 is superimposed.

$$\rho = \frac{\mathrm{cov}(t_1, t_2)}{\sqrt{\mathrm{var}(t_1)\ \mathrm{var}(t_2)}} \qquad (3)$$

where $\mathrm{cov}(t_1, t_2)$ denotes the covariance of the performance of the two algorithms. The effect of cooperation, when manifested in negative correlations, is to increase the performance as well as reduce the risk (Fig. 5). This change is easily understood in that negative correlations mean that one algorithm is particularly good precisely on those cases that are more difficult for the other one, and vice versa. This allows the portfolio, which terminates as soon as the first algorithm completes, to work even better than when the individual algorithms are independent. In the case of the graph-coloring problem, cooperation can be implemented by allowing an algorithm to use incomplete assignments of colors to nodes, posted to a common memory by another algorithm, as a "hint" in its own further explorations (*19*).

This economics approach, emphasizing risk as well as expected performance, has applicability far beyond the solution of NP-complete problems with heuristics, for it addresses any problems that involve variability in performance. In the example of the World Wide Web, one can use a restart strategy where one collects access time statistics, which play the same role as time series in financial markets. The data can then be used to generate performance versus risk curves that specify how to resolve the dilemma of either restarting a request that is taking a long time or waiting in case the Web page will appear in the next few seconds. We tested this scheme by collecting access times for a periodically requested page on the Web. The results show that there are particular periods during the day when the distribution of the access times undergoes qualitative changes. During low congestion periods, the distribution has a relatively small variance in access time, and during high congestion periods, the distribution of access times has a larger variance with an extended tail. Using data from the high congestion period, we varied the time before a restart and found that although expected access time under such a strategy could only be reduced slightly, the standard deviation or risk was reduced by nearly 20%.

Another interesting extension of this methodology is the possibility of dynamically changing the strategy online, so that it can either adapt to changing conditions or optimally exploit information gained in time about an unknown environment. For example, suppose one wishes to use a portfolio strategy with two identical algorithms whose distributions are, however, unknown. This situation is described by Eq. 2 when $p_1(t) = p_2(t)$ is unknown. With a maximum likelihood estimate of $p_1(t)$ (*20*), it is possible to exploit optimally observations of $p(t)$ while dynamically adjusting f_1 in order to get progressively better estimates of $p_1(t)$ and thereby converge on an efficient value of f_1 as more information is received.

An important generalization is provided by the way Monte Carlo algorithms for optimization are usually implemented. If one considers the time it takes to find an acceptable move in an optimization problem, it will have the characteristics of a Las Vegas–type algorithm. We tested this idea by using a simulated annealing algorithm on a vector-quantization problem associated with clustering and found a speed improvement of 5 to 10 times (*20*).

A more exotic example of the applicability of this approach is the construction of a portfolio for a general database search that exploits the properties of quantum computation (*21*). Because the probability distribution of search times in such cases is known beforehand (*22*), the methods presented here can be used to optimize the tradeoff between risk and expected search time.

Given the present trend toward electronic commerce and the explosive use of the Internet, this economics framework can play an important role in allocating computational resources, and thus money, to complete tasks efficiently.

REFERENCES AND NOTES

1. F. Alizadeh, R. M. Karp, D. K. Weisser, G. Zweig, in *Proceedings of the Fifth Annual ACM-SIAM Symposium* (Association for Computing Machinery, New York, 1994), pp. 489–500.
2. M. Mezard, G. Parisi, M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
3. G. Reinelt, *The Traveling Salesman: Computational Solutions for TSP Applications* (Lecture Notes in Computer Science 840, Springer-Verlag, Berlin, 1994).
4. M. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness* (Freeman, San Francisco, 1970).
5. R. Motwani and P. Raghavan, *Randomized Algorithms* (Cambridge Univ. Press, Cambridge, 1995).
6. H. Alt, L. J. Guibas, K. Mehlhorn, R. M. Karp, A. Wigderson, *Tech. Rep. TR-91-057* (International Computer Science Institute, Berkeley, CA, 1991).
7. M. Luby and W. Ertel, *Tech. Rep. TR-93-041* (International Computer Science Institute, Berkeley, CA, 1993).
8. M. Luby, A. Sinclair, D. Zuckerman, *Tech. Rep. TR-93-010* (International Computer Science Institute, Berkeley, CA, 1993).
9. W. Ertel, in *Parallelization in Inference Systems*, B. Fronhofer and G. Wrightson, Eds. (Springer, Dagstuhl, Germany, 1992), pp. 195–209.
10. H. Markowitz, *J. Financ.* **7**, 77 (1952).
11. W. F. Sharpe, *Portfolio Theory and Capital Markets* (McGraw-Hill, New York, 1987).
12. S. Minton, M. D. Johnston, A. B. Philips, P. Laird, *Artif. Intell.* **58**, 161 (1992).
13. D. S. Johnson, C. R. Aragon, L. A. McGeoch, C. Schevon, *Oper. Res.* **39**, 378 (1991).
14. B. Selman, H. Levesque, D. Mitchell, in *Proceedings of the Tenth National Conference on Artificial Intelligence* (AAAI Press, Menlo Park, CA, 1992), pp. 440–446.

15. T. Hogg, B. A. Huberman, C. Williams, *Artif. Intell.* **81**, 1 (1996); S. Kirkpatrick and B. Selman, *Science* **264**, 1297 (1994).
16. T. Hogg and C. P. Williams, in *Proceedings of the Twelfth National Conference on Artificial Intelligence* (AAAI Press, Menlo Park, CA, 1994), pp. 331–336.
17. I. Gent and T. Walsh, *Artif. Intell.* **81**, 59 (1996); B. Selman and S. Kirkpatrick, *ibid.*, p. 273.
18. S. H. Clearwater, B. A. Huberman, T. Hogg, *Science* **254**, 1181 (1991); D. Aldous and U. Vazirani, in *Proceedings of the 35th Symposium on the Foundations of Computer Science*, S. Goldwasser, Ed. (IEEE Press, Los Alamitos, CA, 1994), pp. 492–501; N. Cesa-Bianchi *et al.*, in *Proceedings of the 25th Annual ACM Symposium on the Theory of Computing* (Association for Computing Machinery, New York, 1993), pp. 382–391.
19. T. Hogg and C. P. Williams, in *Proceedings of the Eleventh National Conference on Artificial Intelligence* (AAAI Press, Menlo Park, CA, 1993), pp. 231–236.
20. G. Roumeliotis and B. A. Huberman, in preparation.
21. L. K. Grover, in *Proceedings of the 28th Annual Symposium on the Theory of Computing* (Association for Computing Machinery, New York, 1996), pp. 212–219.
22. M. Boyer, G. Brassard, P. Hoyer, A. Tapp, in *Proceedings of the Workshop on Physics and Computation (PhysComp96)*, in press.
23. R.M.L. acknowledges the support of a Citibank fellowship at Stanford University.

9 August 1996; accepted 28 October 1996

Femtosecond Dynamics of Excited-State Evolution in $[Ru(bpy)_3]^{2+}$

Niels H. Damrauer, Giulio Cerullo,* Alvin Yeh, Thomas R. Boussie, Charles V. Shank, James K. McCusker†

Time-resolved absorption spectroscopy on the femtosecond time scale has been used to monitor the earliest events associated with excited-state relaxation in tris-(2,2′-bipyridine)ruthenium(II). The data reveal dynamics associated with the temporal evolution of the Franck-Condon state to the lowest energy excited state of this molecule. The process is essentially complete in ~300 femtoseconds after the initial excitation. This result is discussed with regard to reformulating long-held notions about excited-state relaxation, as well as its implication for the importance of non-equilibrium excited-state processes in understanding and designing molecular-based electron transfer, artificial photosynthetic, and photovoltaic assemblies in which compounds of this class are currently playing a key role.

Many of the photochemical and photophysical properties of molecules depend upon the kinetics of excited-state processes that occur after the absorption of a photon. Therefore, it is important to understand how excited states behave as a function of time. The conventional view of this temporal evolution holds that photoreactivity is largely dictated by the characteristics of the lowest energy excited state of a molecule. Thus, higher energy excited states are presumed to convert to this lowest energy state and in so doing are removed from any functional role in photochemical and photophysical transformations. Femtosecond time-resolved spectroscopy (*1*) has resulted in experimental observations that call into question the validity of this model; striking examples include the 200-fs *cis*-to-*trans* isomerization of rhodopsin (*2*), rapid photodissociation of CO from myoglobin-CO (*3*), and ultrafast electron injection into dye-sensitized semiconductor electrodes (*4*). These cases among others reveal a pattern of photoreactivity arising from nonthermalized excited states in which structural rearrangement and electron transfer can kinetically compete with processes such as intramolecular vibrational relaxation (IVR), internal conversion (IC), and intersystem crossing (ISC).

The inference that nonequilibrated excited states can play a chemically significant role in photoinduced transformations could have important consequences in a variety of areas ranging from design principles for electron-transfer assemblies and photochemical energy storage devices to the formulation of new theoretical models for molecular-based energy conversion and excited-state relaxation dynamics. Although much of the work in the ultrafast dynamics community has concentrated on either small molecules or biological systems, our research focuses on transition metal compounds (*5*, *6*). Considerable effort is being expended in many laboratories to incorporate such complexes into schemes for artificial photosynthesis (*7*), photocatalysis (*8*), and the development of molecular-based photovoltaic and opto-electronic devices (*9*). In addition, the importance of ISC and IC processes in the photoinduced properties of metal-containing complexes makes such systems of interest for ultrafast dynamical studies of their excited-state behavior (*10*). We have obtained results that are not consistent with conventional models for describing photoinduced dynamics in transition metal complexes, suggesting the need to reevaluate currently accepted views of their excited-state behavior.

Tris-(2,2′-bipyridine)ruthenium(II), or $[Ru(bpy)_3]^{2+}$,

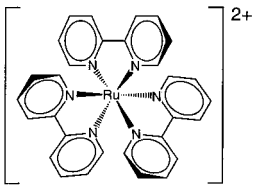

is representative of a class of molecules that has played a central role in the development of inorganic photophysics in addition to providing the underpinning for the last two decades of research on transition metal–based photosensitization, charge separation, and photoinduced electron transfer chemistry (*11*). We have therefore chosen it as a prototype for our study of the ultrafast dynamics of metal complexes. The strong visible absorption characteristic of this molecule (Fig. 1) can be described as a metal-to-ligand charge transfer ($^1MLCT \leftarrow {}^1A_1$), in which an electron located in a metal-based *d*-orbital is transferred to a π^* orbital of one of the bpy ligands ($h\nu$, photon energy) (*12*). The excited-state species that is eventually

$$[Ru^{II}(bpy)_3]^{2+} \xrightarrow{h\nu} [Ru^{III}(bpy^-)(bpy)_2]^{2+}$$

formed (a ^3MLCT state) is well known to engage in both oxidative and reductive chemistry (*11*). This capability, coupled with its relatively long lifetime in fluid solution ($\tau \approx 1$ μs), near unity quantum yield of formation (*13*), the high visible absorptive cross section of the ground state, and the overall photochemical stability of this molecule and its derivatives makes them amenable to a wide variety of applications (*14*, *15*). We have used femtosecond absorption spectroscopy to time resolve the formation of the ^3MLCT state in $[Ru(bpy)_3]^{2+}$ (*16*) and have observed the initial evolution of the Franck-Condon state.

The laser system used has been described in detail elsewhere (*17*, *18*). Excited-state difference spectra at various time delays Δt (Fig. 2) show that spectral changes in the 450- to 490-nm range are quite dramatic: A bleach begins to evolve at $\lambda = 470$ nm near

N. H. Damrauer, T. R. Boussie, J. K. McCusker, Department of Chemistry, University of California, Berkeley, CA 94720, USA.
G. Cerullo and A. Yeh, Material Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA.
C. V. Shank, Department of Chemistry, University of California, Berkeley, CA 94720, and Materials Science Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA.
*Present address: Dipartimento di Fisica del Politecnico, P.za L. Da Vinci 32, 20133 Milano, Italy.
†To whom correspondence should be addressed.