



Performance Variability and Project Dynamics

BERNARDO A. HUBERMAN
HP Labs, 1501 Page Mill Rd., Palo Alto, CA 94304
email: bernardo.huberman@hp.com

DENNIS M. WILKINSON
HP Labs, 1501 Page Mill Rd., Palo Alto, CA 94304 and Stanford University Physics Department, 382 Via Pueblo Mall, Stanford, CA 94305
email: dennis.wilkinson@hp.com

Abstract

We present a dynamical model of complex cooperative projects such as large engineering design or software development efforts, comprised of concurrent and interrelated tasks. The model contains a stochastic component to account for temporal fluctuations both in task performance and in the interactions between related tasks. We show that as the system size increases, so does the average completion time. Also, for fixed system size, the dynamics of individual project realizations can exhibit large deviations from the average when fluctuations increase past a threshold, causing long delays in completion times. These effects are in agreement with empirical observation. We also show that the negative effects of both large groups and long delays caused by fluctuations may be mitigated by arranging projects in a hierarchical or modular structure. Our model is applicable to any arrangement of interdependent tasks, providing an analytical prediction for the average completion time as well as a numerical threshold for the fluctuation strength beyond which long delays are likely. In conjunction with previous modeling techniques, it thus provides managers with a predictive tool to be used in the design of a project's architecture.

Keywords: team performance, team dynamics, large cooperative projects, stochastic dynamics

1. Introduction

One of the main challenges in a large design project or, more generally, any problem solving process is the coordination of the different tasks comprising that project or process (Galbraith, 1977; Thompson, 1978; Hogg and Huberman, 1993). Coordination is particularly difficult in situations where the complexity of the project leads to its division into concurrent, interdependent tasks whose results must be dynamically integrated into an overall satisfactory solution. Examples are large-scale software design projects, engineering design projects, and industrial research and development efforts. It has long been recognized that because of the coupled nature of the component tasks such problem solving processes are inherently iterative in their execution (Kline, 1985; Whitney, 1990). Information from the partial solution to a given task can trigger a chain of revisions as solutions to other, related tasks are modified (Allen, 1966; Clark and Fujimoto, 1991; Thomke, 1997; van Zandt, 1999).

A number of formal models of cooperative processes have been proposed to predict the number of iterations required for completion, or to suggest optimal concurrency and iteration schemes which minimize the overall time or cost (e.g., Ahmadi and Wang (1994);

Ha and Porteus (1995); Krishnan et al. (1997); Jin and Levitt (1996); Ford and Sterman (1998); Roemer et al. (2000); Loch et al. (2001); Peña-Mora and Li (2002)). One example of these approaches is the work transformation matrix (WTM) model (Smith and Eppinger, 1997), which is an extension of the design structure matrix model (see Browning (2001) for a review). It provides a simple mathematical representation of design iterations that estimates the number of iterations and work required for a given arrangement of coupled tasks.

Helpful as these models are, they do not account for the frequent situation in which projects enter a vicious cycle of continuing revisions (Cusumano and Selby, 1995; Joglekar, 2001; Terwiesch et al., 2002) resulting in budget overruns, missed business opportunities and delays which, at times, can make the final solution obsolete (Brooks, 1975; Morris and Hugh, 1987; Terwiesch and Loch, 1999). A major factor leading to these undesirable outcomes stems from the unpredictable fluctuations in the value that a particular unit of work on a single task brings to the overall project. In order to take into account these fluctuations, recent models have focused on individual sources of variability, including asynchronicity and random timing of task updates and information exchange (Mihm et al., 2003; Huberman and Glance, 1993), information withholding (Yassine et al., 2003), exogenous changes (Tamai and Itou, 1997; Mar, 1999), volatility of resource allocation (Repenning et al., 2001), behavioral choice (Ford and Sterman, 1999), uncertainty of performance evaluation (Browning et al., 2002), and the complicated landscape of performance maxima and minima (Mihm et al., 2003). However, since these models are not stochastic in nature, they do not account for the random and unpredictable nature of project dynamics, which is an intrinsic and unavoidable element of any complex cooperative project.

In Carrascosa et al. (1998), a stochastic effect is introduced into a model of the evolution of an engineering design project, and its effect on the probability of project completion within a given time is studied for one particular project using simulation. Similarly, Browning and Eppinger (2002) includes a stochastic component of a very complex nature in the study of different architectures for a given project, and Ahmadi et al. (2001) include two types of Markov learning processes in a framework used to determine optimal structures for minimizing design iterations. These models provide tools for predicting project evolution in a given situation by means of simulation. In contrast, the goal of this paper is not only to provide a tool for prediction, but also to make a general study of the effect of project size, interaction strengths, and stochastic fluctuations on project dynamics.

As we show, the evolution of a set of coupled tasks can vary greatly depending on the structure of the interdependencies, fluctuation strength, and other parameters, so that a general treatment is necessary to study these effects. The simple formulation of our model and the analytical results it provides clearly elucidate the role played by these parameters in a project's evolution.

In this paper, we thus present a general study of iterative problem solving processes using a mathematical model which explicitly incorporates the fluctuating nature of task performance and interdependence. The model uses the well-established formalism of the work transformation matrix (WTM) to describe the dynamics of the process. In the WTM model, progress on a given project is described by a state vector which represents how much work is left to be done on each of the component tasks. The work rates and interactions between tasks are encapsulated in the work transformation matrix itself. As the project

proceeds, the state vector is updated at each time step according to multiplication by the WTM.

The crucial fluctuating component is included by varying the elements of the WTM at each updating step. Progress towards an overall solution is thereby represented by a stochastic dynamic process. Such a process captures the erratic nature of the evolution of real-world projects and their component tasks. Furthermore, as we discuss below and show in this paper, the mathematical effect of fluctuations and thus the behavior predicted by our model is consistent with previous empirical results on complex projects, as well as our own and other authors' computer simulations.

Moreover, our analytical results are general and may be applied as a predictive tool to any given project, using numerical inputs relevant to the project which may be obtained as described in Browning (2001) and particularly McDaniel (1996) (see also a brief summary of this method in Section 2 of this paper). The dynamics of a project consisting of even four or five interrelated tasks can be complex and difficult to understand without the use of a mathematical model (see e.g. Smith and Eppinger (1997)), and the inclusion of fluctuations further complicates matters because their effect is highly system-dependent. Our model, which is simple to implement once the inputs are determined, may therefore be of interest as a tool for understanding and predicting the behavior of a given project.

In particular, our results are as follows. We first show that the time to solution increases on average as the number of interactions in the project increase, in agreement with empirical results (Clark, 1989; Griffin, 1997; Reel, 1999; Tamai and Itou, 1997) and previous simulations (Mihm et al., 2003; Yassine et al., 2003). This property of the WTM model is independent of the fluctuations but has not been definitively presented in previous work. We then show that as the average strength of the fluctuations increases, the time to completion increases, in agreement with empirical studies (Ibbs, 1997). We demonstrate that in a large project comprising many interactive tasks, a hierarchical or modular structure can, on average, alleviate the problem of large convergence time, as previously proposed (e.g. Hammer (1996)) and in agreement with empirical results (Ulrich, 1995; Sosa, 2003) and previous studies (Mihm et al., preprint; Rivken and Siggelkow, 2003; Ethiraj and Levinthal, 2004).

Moreover, because of the temporal variability of the fluctuations, evolution towards solution of any particular instance of a project can differ greatly from the average behavior. As a result, projects which would converge smoothly to solution in the absence of fluctuations can deviate significantly from this path. When the temporal variability is low, convergence to a solution is smooth and the finishing times are approximately normally distributed close to the average value. But above a given threshold in variability, the distribution of finishing times undergoes a transition to a heavy-tailed log-normal form. This distribution is in agreement with earlier empirical results on cooperative problem solving (Clearwater et al., 1991) and implies possible finishing times far greater than the average. We provide analytical results which may be applied, using the inputs of a given project, to predict the average time to solution as well as the likelihood of significant deviations above this average. A critical value for the strength of the fluctuations, above which fluctuations can cause significant delays, is also provided.

Finally, we show that the effect of fluctuations is more severe in projects which converge slowly to solution. This implies that hierarchical organization not only can decrease the

average time to solution, but can also mitigate the possibly unavoidable effects of fluctuating interactions.

The paper is organized as follows. In Section 2 we present and discuss the dynamical model. In Section 3 we study the model's average behavior and the relation between convergence time and project size. In Section 4 we consider the effects of fluctuating efforts and establish regimes that lead to large deviations away from the smooth convergence to the goal. Section 5 summarizes our results.

2. A Model of Group Problem Solving

Consider a group of individuals, or even computer programs, working cooperatively towards the solution of a problem. An example of this process in the design space is provided by a large software development effort in which individuals or teams work iteratively on pieces of the software and pass it along to others so that whole modules and eventually the whole application may be completed. Alternatively, one can conceive of teams of engineers involved in the design of a complicated mechanism that requires tight integration of all the parts in order to achieve the desired goal.

The progress of a cooperative project comprising n tasks towards a solution can be represented mathematically by a state vector \mathbf{x} with n components, as in the WTM model of Smith and Eppinger (1997). The i th element of \mathbf{x} represents how far task i is from completion. The dynamics of the system are specified by a time-dependent interaction matrix \mathbf{A}_t which encapsulates the interdependencies among the tasks. The state of the overall project at time $t + 1$ is determined by

$$\mathbf{x}_{t+1} = \mathbf{A}_t \mathbf{x}_t, \quad (1)$$

where \mathbf{x}_t is the state at time t . The initial state \mathbf{x}_1 may be taken to be a vector of 1's, defining the scale by which we measure how much work is left on the tasks. Ideally, as the project proceeds, the entries of the state vector will become smaller and smaller as the tasks near completion.

The interpretation of the elements of the interaction matrix is as follows (see also the example below). The off-diagonal elements measure how much one task's partial solution helps or hinders the completion of the other tasks. A positive entry signifies that one unit of work on task j causes $(A_t)_{ij}$ units of rework on task i , while a zero entry means that tasks i and j have no direct effect on each other. In the WTM model, only positive or zero entries were allowed. Our model allows for negative entries when efforts on task j 's hasten the completion of task i . The diagonal elements of the interaction matrix \mathbf{A}_t account for different rates of progress on different tasks¹. This is in contrast to the WTM model where all tasks are performed at the same rate.

To illustrate how the WTM model works, consider as an example the following hypothetical development of an internet-based software client. Please note that the WTM is an established tool and matrices such as the one below have been calculated and applied in real-world settings (e. g. McDaniel, 1996), as further discussed below; this example was chosen for simplicity. The software development process has been divided into three

modules: user interface, database front end, and network layering. It is thus modeled by a three-dimensional system as follows:

$$\mathbf{x}_t = \begin{pmatrix} \text{work remaining on user interface} \\ \text{work remaining on database front end} \\ \text{work remaining on network layering} \end{pmatrix}.$$

The initial state \mathbf{x}_1 of the system is a vector of 1's and the time step is one week. After a week, the state of the system is given by

$$\mathbf{x}_2 = \mathbf{A}_1 \mathbf{x}_1 = \begin{pmatrix} 0.65 & 0.2 & 0.05 \\ 0.3 & 0.6 & 0.07 \\ 0.09 & 0.05 & 0.7 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 0.9 \\ 0.97 \\ 0.84 \end{pmatrix}, \quad (2)$$

where \mathbf{A}_1 is a hypothetical interaction matrix which describes the dynamics over the first week of the project. To further illustrate how the model works, consider in particular the 0.65 entry in the (1, 1) position of \mathbf{A}_1 . When the state vector is updated, this diagonal element multiplies the work remaining on task 1. Its value thus means that, were there no task interdependency, 65% of the work remaining on task 1 would be completed this week. However, progress on other tasks may create more work for task 1, as modeled by off-diagonal elements such as the 0.05 in the (1, 3) position of \mathbf{A}_1 . This number indicates that work on task 3 adds to the work on task 1 by a factor of 5%. In this particular example, all the off-diagonal elements are positive but in general they need not be; see the section on fluctuations, below, for an example and discussion.

Notice how in this example the interdependencies between tasks slow down the overall progress. The tasks proceed rather quickly when taken alone (diagonal elements), but due to the strong interdependencies (off-diagonal elements), not much progress is made overall in the first week. These off-diagonal elements were chosen to reflect the difficulties in making components of a large software effort compatible.

2.1. Project Architecture

The matrices \mathbf{A}_t entirely define the dynamics of the system by representing the interactions between tasks (off-diagonal elements) and the task work rates (diagonal elements). To a large degree, properties of the project's architecture define the elements of the interaction matrices. Such properties include the inherent interdependence of related tasks, the structure of the organization (hierarchical, flat, etc.), and the relationships between individuals working on tasks and communities of practice (Huberman and Hogg, 1995), among others.

Interdependencies of this type give a constant, underlying structure to the system's dynamics and form a basic interaction matrix \mathbf{A}_0 . This matrix is simply the work transformation matrix (Smith and Eppinger, 1997), which we will refer to as the *unperturbed* interaction matrix or unperturbed WTM. The actual interaction matrices \mathbf{A}_t in our model are created

by introducing fluctuations into the elements of the unperturbed WTM. In practice, an unperturbed WTM can be created for a particular project by assigning numerical values to the project's design structure matrix (DSM). The DSM is an important tool (see Browning, 2001 for a review) which encapsulates a project's interdependencies, usually by means of a survey of the individuals involved. For example, McDaniel (1996) created a WTM for the appearance design of a vehicle's interior by assigning each qualitative DSM entry a numerical value and averaging the values obtained from many sources. This process is shown in Tables 1 and 2. The numerical values chosen for the interdependencies are of crucial importance to the system's evolution, as explained in Sections 3 and 4.

Table 1. Sample DSM characterizing interdependencies in appearance design of vehicle interior. Here S means strong, M medium, W weak and 0 none.

	1	2	3	4	5	6	7	8	9	10
1. Carpet		W	0	W	W	0	0	0	W	0
2. Center console	W		W	0	0	S	W	0	M	0
3. Door trim panel	0	W		W	0	M	0	0	M	W
4. Garnish trim	W	0	M		0	W	W	0	W	0
5. Overhead system	W	0	0	0		0	0	0	0	0
6. Instrument panel	0	S	S	W	0		W	0	0	M
7. Luggage trim	0	0	0	W	0	W		W	W	0
8. Package tray	0	0	0	W	0	0	W		M	0
9. Seats	W	M	M	W	0	W	W	M		M
10. Steering wheel	0	0	0	0	0	S	0	0	M	

Table 2. WTM for vehicle interior appearance design process (McDaniel, 1996) created from averaging many DSM's like the one shown in Table 1. Here S interdependencies were given the value 0.15, M 0.1, and W 0.05. The diagonal elements were obtained by estimating the autonomous completion time for each component.

	1	2	3	4	5	6	7	8	9	10
1. Carpet	0.85	0.06	0.01	0.03	0.03	0	0	0	0.03	0
2. Center console	0.05	0.53	0.02	0	0	0.15	0.01	0	0.12	0.01
3. Door trim panel	0.01	0.02	0.47	0.04	0	0.12	0.01	0	0.09	0.01
4. Garnish trim	0.03	0	0.09	0.68	0	0.07	0.05	0.01	0.04	0
5. Overhead system	0.02	0	0	0	0.83	0	0	0	0	0
6. Instrument panel	0	0.15	0.13	0.08	0	0.28	0.03	0	0.01	0.1
7. Luggage trim	0	0.01	0.01	0.05	0	0.03	0.76	0.03	0.04	0
8. Package tray	0	0	0	0.05	0	0	0.03	0.83	0.08	0
9. Seats	0.04	0.12	0.09	0.04	0	0.02	0.02	0.08	0.63	0.1
10. Steering wheel	0	0.01	0.01	0	0	0.13	0	0	0.1	0.7

2.2. *Projects with two or more Different WTM*s

In this paper we consider only time periods for which the unperturbed WTM of a project remains unchanged. The purpose of this caveat is to only to simplify the discussion, as it does not in any way restrict the analysis of this paper. There is evidence, for example, that the interaction strengths in many projects undergo large-scale change between the initial, middle and final phases of a project (Morris, 1982). Alternatively, a project may undergo a major reorganization. These situations would be accounted for by using separate WTM for each phase; the analysis presented in this paper would then apply separately to each phase of the dynamics.

2.3. *Errors in the WTM*

For practical applications, it is of interest to understand the effect of an error in the unperturbed WTM. To our knowledge, there has been no previous discussion of such errors in the context of the WTM formalism. Nevertheless, it may be difficult in practice to obtain accurate values for all WTM entries in a particular application.

The effect of errors in an unperturbed WTM is highly system-dependent and varies with the size of the errors, the number of errors, and the location of the errors in the matrix in a complex way. For most systems, the WTM method is tractable if the errors are small, mostly because the elements of unperturbed WTM are almost all positive. Of course, care must be taken because in certain positive matrices small errors can have a significant effect on system dynamics. This subject is beyond the scope of this paper; for more information see Golub and van Loan (1989) or Mine (1998).

In our model, fluctuations are introduced into the WTM as described below. This can lessen the severity of errors because, as we will see, our model predicts range of probable outcomes. This is as opposed to previous methods which predict only one outcome. Even if there are errors in the unperturbed WTM, the actual behavior may still lie within the range predicted by our model. Moreover, if a manager feels particularly uncertain about a certain entry of the WTM, he can simply assign a greater fluctuation value to this entry.

Another strength of our model is its ability to identify those systems in which small errors in the WTM entries may cause a large errors in the predictions. This is because the severity of errors in the WTM is strongly related to the severity of the effect of fluctuations. As we will see, the effect of fluctuations increases as the parameter ε , as defined in equation 13, increases. A large value of ε , for example on the order of 0.25 or more, thus indicates that errors in the WTM will have a particularly severe effect and extra care must be taken in such systems to assure the accuracy of the WTM elements.

2.4. *Fluctuations*

Interactions in a cooperative problem solving process do not remain constant on short time scales but vary due to a number of features, including asynchronicity of information exchange, exogenous changes, behavioral choice, uncertainty of performance evaluation, the varying relevance of the partial solutions of related tasks to each other, and other

factors as discussed in the introduction. These factors conspire to produce fluctuations in the elements of the interaction matrix. The exact nature of the fluctuations in any project is impossible to determine in advance because of the complexity of the process and of these interacting factors. The problem solving effort is thus effectively a stochastic process, where the fluctuations introduce a random component into the dynamics.

To see how these stochastic fluctuations affect the dynamics, consider the software development example introduced above. The evolution of the project in the first week was determined as shown in Eq. (2). During the second week, a different story unfolds, as the fluctuations alter the entries of the interaction matrix :

$$\mathbf{x}_3 = \mathbf{A}_2 \mathbf{x}_2 = \begin{pmatrix} 0.8 & 0.27 & -0.05 \\ 0.23 & 0.55 & 0.04 \\ -0.03 & 0.02 & 0.7 \end{pmatrix} \begin{pmatrix} 0.9 \\ 0.97 \\ 0.8 \end{pmatrix} = \begin{pmatrix} 0.9419 \\ 0.7725 \\ 0.5524 \end{pmatrix}.$$

Note in particular that the interaction matrix has changed, so that the state vector is updated in a different way from the way it changed after the first week. During the second week, progress was slow on the user interface component (interaction matrix element (1,1)). Moreover, hindering interactions, perhaps due to a lack of communication with the front end developer (elements (1,2) and (2,1)), resulted in lost ground on this task. The front end itself was hurt by the bad communication but programmer 2 worked hard (element (2,2)) and made progress. The network protocol task also made good progress as its results were fortuitously compatible with those of task 1 (elements (1,3), (3,1)).

Fluctuations are reflected in our model by having the elements of the interaction matrix vary randomly in time with a given distribution. For simplicity, we assume that the distribution does not vary over the lifetime of the project. It is straightforward to relax this assumption if the project undergoes a major reorganization as explained above.

To focus on the effect of the fluctuations, we introduce new notation for the interaction matrices \mathbf{A}_t . Recall that the off-diagonal elements of \mathbf{A}_t indicate how much one unit of work on a certain task helps or hinders a different task at time step t , while the diagonal elements represent the work rate on individual tasks. Consider a particular element of \mathbf{A}_t . One contribution to this element comes from the unperturbed work transformation matrix (recall the abbreviation WTM), which represents some *a priori* prediction of the value of the interaction. The other contribution is from the unforeseen fluctuations. This split is summarized by:

$$\begin{array}{l} \text{each element of} \\ \text{interaction matrix} \end{array} = \begin{array}{l} \text{unperturbed WTM} \\ \text{(constant matrix)} \end{array} + \begin{array}{l} \text{fluctuation part} \\ \text{(matrix which varies in time)} \end{array}.$$

The fluctuating part of the interaction matrix is itself a matrix of random variables that vary in time with certain means and variances. For each element, the fluctuation mean indicates how far the value of the interaction is on average from the unperturbed WTM value, and the variance indicates how much temporal deviation there is from this mean. Since the fluctuation means are constant throughout the project's lifetime, we can further split the

fluctuation part as follows:

$$\begin{array}{l} \text{each element of} \\ \text{interaction matrix} \end{array} = \begin{array}{l} \text{unperturbed WTM} \\ \text{(constant matrix)} \end{array} + \begin{array}{l} \text{fluctuation mean} \\ \text{(constant matrix)} \end{array} + \\ + \begin{array}{l} \text{varying part of fluctuation} \\ \text{(matrix of mean-zero random variables)} \end{array} .$$

The mathematical representation of this relation is

$$\mathbf{A}_t = \mathbf{A}_0 + \mathbf{M} + \mathbf{S}_t, \quad (3)$$

where \mathbf{A}_0 is the unperturbed work transformation matrix, \mathbf{M} is the matrix of mean values of the fluctuations, and \mathbf{S}_t is a time-dependent matrix of mean-zero random variables representing the temporal variation of the fluctuations. Using this notation the dynamics of the overall problem solving process are expressed by the stochastic equation

$$\mathbf{x}_{t+1} = (\mathbf{A}_0 + \mathbf{M} + \mathbf{S}_t)\mathbf{x}_t. \quad (4)$$

The state vector \mathbf{x} is a vector of random variables, and the dynamics of the problem solving process are characterized by the system's average value $\langle \mathbf{x}_t \rangle$ and its moments $\langle |\mathbf{x}_t|^p \rangle$. The moments provide a measure of how likely the system is to be found far from its average behavior.

2.5. Tasks are Performed in Parallel

In this paper we treat only coupled tasks which must be performed in parallel. That is, no task or subgroup of tasks may be performed independently of the others, because it requires input from one or more of the other tasks. This condition in no way restricts the applicability of the analysis in this paper. In projects which may be decomposed into subgroups of tasks that are performed in series, our analysis would simply apply to each subgroup in the sequence.

2.6. Stopping Criterion

A final important feature of our model is the stopping criterion and the interpretation of the zero vector. We consider the zero vector to represent an "optimal" solution, and the values of the state vector as an abstract measure of the amount of work left to be done before a task's solution is optimal. In the evolution described by Eq. (1), even when \mathbf{A}_t is constant and the dynamics are convergent, the zero vector is never reached, as the state vector asymptotically approaches zero as shown in figure 1. In a given project, the goal is thus reached when managers decide that the tasks' values are close enough to zero. This is in contrast to the WTM formulation in which the total work is calculated as the amount required to bring the state vector to a vector of zeros.




Figure 1. Simulation of project dynamics showing the evolution of the first component of the state vector. This plot was created using only the unperturbed interaction matrix A_0 shown in Table 2 (no fluctuations). At left, evolution over 100 weeks; at right, a plot of weeks 60–100 showing the difference in completion time for two different stopping criteria (dotted lines).

Our model thus implicitly contains a stopping criterion that determines when the work is finished. For example, the stopping criterion may be that all the elements of x are less than 0.05. Other, task-specific criteria are of course possible. As shown in the right panel of figure 1, a slight change in the stopping criterion can mean a difference of weeks or months in the completion time. This interpretation is quite reasonable when one considers that real-world projects are always terminated by managers when the solution is deemed satisfactory by some subjective indicators, or when time-to-market or other outside pressure is encountered.

2.7. Continuous Time Model

The model proposed above is discrete in time, implying that all the task states are updated at once. This is applicable to situations where individuals update their information synchronously, such as at a weekly meeting. Conversely, in situations where information is continuously passed between various individuals or subtasks in asynchronous fashion, a continuous-time model is needed (Huberman and Glance, 1993). For an unperturbed system with no fluctuations, the qualitative conclusions are the same for the continuous and discrete model. This is shown in Appendix D where the mathematics of an unperturbed continuous model, mirroring the discrete treatment of the main body of this paper, are presented. Unfortunately, stochastic fluctuations are difficult to treat analytically in the continuous case and we leave this as an area for future work.

2.8. Model Assumptions

The introduction of fluctuations into the interaction matrix means that our model requires few assumptions about the nature of the project. The one assumption that is necessary, just as

in the original WTM treatment, is that the project's evolution can be described using a linear model. A linear model first means that the average work rate on each task remains constant over the project's lifetime. It also means that the effect of one task on another, in relation to the amount of work remaining on that task, remains constant over the project's lifetime. As mentioned above, there is evidence that task interrelations may change as the project enters different phases; as also mentioned above, however, the model can be equally well applied in these cases by using a different unperturbed WTM for each phase and modeling the dynamics separately for the each phases.

3. Average Dynamics and Project Size

Equations (4) and (19) describe the evolution of the cooperative solving process towards a solution. As discussed above, these equations are stochastic and describe a range of possible dynamics. In order to elucidate the possible behaviors, we first study the average dynamics of the process. As we show later, when the fluctuations in interaction strength are small, the evolution of the problem solving process is very likely to be close to the average. For larger noise, individual realizations of a given project may differ greatly from the average. Nevertheless, this quantity is important in studying the dependence on project organization, system size and interaction strengths.

As shown in appendix A, the matrices \mathbf{S}_t do not affect the average dynamics, so the average state of the process (4) evolves according to

$$\langle \mathbf{x}_{t+1} \rangle = (\mathbf{A}_0 + \mathbf{M}) \langle \mathbf{x}_t \rangle \quad (5)$$

This is a simple linear system whose dynamics are well-understood. The convergence properties are completely determined by the magnitude of the largest eigenvalue λ of the average interaction matrix $\mathbf{A}_0 + \mathbf{M}$. That is, if $\lambda < 1$, the process will produce a satisfactory solution as the state vector \mathbf{x} converges asymptotically to the zero vector; but if $\lambda > 1$, the process will not yield a solution as the state vector diverges, or continues to increase. This behavior is demonstrated in figure 2.

Moreover, the value of λ determines the rate of convergence (or divergence), with higher values of λ implying slower convergence (or faster divergence). This property is explained below and is demonstrated in figure 2. This is the reason why the numerical values chosen to represent the qualitative entries of the DSM are so important, as demonstrated in the figure.

The above convergence properties are a fundamental result of linear algebra (Strang, 1976) and have been noted in connection with the WTM model (Smith and Eppinger, 1997). However, previous WTM treatments have not fully explored the elegant simplicity of this linear system, for which the asymptotic solution can be concisely expressed. To do this, we will assume that $\mathbf{A}_0 + \mathbf{M}$ has a simple, dominant eigenvalue². Doing so we obtain (see Appendix B for a derivation) the following simple expression for the long-time solution to Eqs. (5):

$$\langle \mathbf{x}_t \rangle = c \mathbf{u} \lambda^t \quad (6)$$




Figure 2. Simulation of average dynamics of processes with discrete dynamics and various values of λ . This plot was created using the DSM interaction matrix for the appearance design of a vehicle interior (Table 1) by simply varying the numerical values assigned for strong, medium, and weak dependencies. The diagonal entries reflect the autonomous task completion rate and are those of the WTM in Table 2. Here $\mathbf{M} = \mathbf{0}$.

where the constant c is given by $c = \mathbf{v} \cdot \mathbf{x}_0$,³ and \mathbf{u} and \mathbf{v} are the right and left eigenvectors, respectively, of the matrix $\mathbf{A}_0 + \mathbf{M}$ corresponding to the eigenvalue λ . Expressions (6) clearly show that the convergence rate is determined by λ . It is important to note that the long time limit may have different meanings for different systems. If the second largest eigenvalue is very small compared to λ , the above asymptotic limit may be reached very quickly, e.g. after as few as 2 or 3 iterations.

The dependence of the largest eigenvalue on matrix properties relates the average convergence rate to parameters such as system size and organizational structure. If the organization of the problem solving process is flat, i.e., the level of interaction between all tasks is arbitrary, most of the entries of the interaction matrix are non-zero. In this case, the largest eigenvalue of $\mathbf{A}_0 + \mathbf{M}$ is given on average by Juhász (1982)

$$\lambda \sim n\mu + O(\sqrt{n}), \quad (7)$$

where n is the size of the matrix $\mathbf{A}_0 + \mathbf{M}$ and μ is the average value of its elements. If we further assume that the interactions are symmetric, that is, individual i 's results have the same effect on individual j that individual j 's have on individual i , a more precise result is available (Füredi and Komlós, 1981):

$$\lambda = n\mu + \sigma^2/\mu, \quad (8)$$

where σ^2 is the variance of the elements of $\mathbf{A}_0 + \mathbf{M}$. The linear growth of λ with the number of independent tasks (n) and the tendency of interactions to hinder progress (μ) shows that the time to solution for a flat organizational structure grows quickly as the




Figure 3. Largest eigenvalue of matrices of different sizes. Samples of 1000 matrices for each $n = 3$ through 15 were generated by choosing the matrix elements randomly from a distribution with average $\mu = 0.15$ and standard deviation 0.3. The solid line connects the average value of λ for each sample, and the error bars are the standard deviations. The dotted line shows the analytic prediction $\langle \lambda \rangle \sim n\mu$.

project complexity increases. The growth of λ with n is demonstrated numerically in figure 3.

In many situations, however, smaller teams work on parts of a problem which are then put together at a higher organizational level. From the point of view of our dynamical model, this means that each task interacts only with the other tasks of the team or their higher order aggregates. The ensuing interaction matrix will then have blocks of nonzero elements and zero entries elsewhere. In this case, a slower growth of the largest eigenvalue with size is obtained (Hogg et al., 1989):

$$\lambda \sim \mu \ln n, \quad (9)$$

implying that much larger groups can be stable when they are structured in highly clustered fashion, i.e. teams.

4. The Effect of Fluctuations

We now examine the effect of fluctuations on the dynamics of cooperative problem solving. As discussed above, these fluctuations are a key ingredient of any cooperative process and are due to many different factors. Recall that the fluctuations are included in the model by modifying the interaction matrix from the simple unperturbed WTM to the matrix $\mathbf{A}_t = \mathbf{A}_0 + \mathbf{M} + \mathbf{S}_t$, where \mathbf{M} is the matrix of the mean values of the fluctuations and \mathbf{S}_t is a matrix of mean-zero random variables representing their temporal variation.

In the absence of fluctuations, the process is governed entirely by the unperturbed matrix \mathbf{A}_0 . The dynamics are thus equivalent to the average dynamics described in Section 3,

with the interaction matrix \mathbf{A}_0 instead of $\mathbf{A}_0 + \mathbf{M}$. That is, the process either converges smoothly to solution, if the largest eigenvalue of \mathbf{A}_0 is less than 1, or diverges smoothly away from solution if this eigenvalue is greater than 1. Moreover, for convergent processes, the magnitude of the largest eigenvalue of \mathbf{A}_0 determines the time to solution; processes with larger eigenvalues take longer to converge. This behavior was previously discussed and demonstrated in figure 2, above.

The fluctuations affect the dynamics of the process in two distinct ways. First, if the fluctuations have a non-zero mean, that is, they tend to help or hinder task completion, \mathbf{M} is non-zero and the average dynamics are different from the unperturbed dynamics. Second, if the strength of the temporal variation S_t of the fluctuations exceeds a certain threshold, the dynamics of individual project realizations may deviate substantially from the smooth evolution of the average dynamics, possibly causing long delays. These two effects are explored in Sections 4.1 and 4.2.

4.1. Shifted Average Dynamics

As mentioned above, average values of the fluctuations have the effect of shifting the average dynamics away from the unperturbed dynamics.

The unperturbed system is given by $\mathbf{x}_{t+1} = \mathbf{A}_0 \mathbf{x}_t$, and the convergence rate is thus determined by the largest eigenvalue λ_0 of \mathbf{A}_0 . When fluctuations are added, the average evolves according to

$$\langle \mathbf{x}_{t+1} \rangle = (\mathbf{A}_0 + \mathbf{M}) \langle \mathbf{x}_t \rangle,$$

and the average convergence rate is shifted from λ_0 to λ , the largest eigenvalue of $\mathbf{A}_0 + \mathbf{M}$. In this section for simplicity we will proceed using the discrete time formulation only; the results which follow are applicable in the continuous model as well.

Myriad forms are possible for the average fluctuation matrix \mathbf{M} . A reasonable form to consider is that where each element of \mathbf{M} is proportional to the corresponding element of \mathbf{A}_0 by some factor m . That is, if the (2, 3) element of \mathbf{A}_0 equals 0.4, indicating a strongly hindering correlation, we would have $M_{23} = 0.4m$, while for a weakly hindering element such as $(A_0)_{12} = 0.03$ we would take $M_{12} = 0.03m$. This assumes that the relative strength of the correlations is not affected on average by the fluctuations. One can thus model a wide variety of behavior by varying m , even to have negative values. However, the most reasonable values for m are positive, since unforeseen events typically hinder rather than help the completion of the project on average.

The average interaction matrix thus takes the form

$$\mathbf{A}_0 + \mathbf{M} = (1 + m)\mathbf{A}_0,$$

and its largest eigenvalue is given simply by

$$\lambda = (1 + m)\lambda_0.$$

This value determines the average dynamics of the perturbed system, as described in Section 3. In particular, for positive values of m , the average effect of the fluctuations is to increase the time to solution. Moreover, if $\lambda_0 < 1$ but

$$m > \frac{1}{\lambda_0} - 1,$$

the fluctuations cause the project to diverge on average even though an unperturbed treatment would predict convergence. This scenario is more likely the closer λ_0 is to 1, i.e., the slower the predicted convergence in the unperturbed system.

For an example, consider the average dynamics shown in figure 2, above. The curves in this plot were generated by varying the off-diagonal entries of \mathbf{A}_0 from Table 2, above, taking $\mathbf{M} = \mathbf{0}$. Note that a similar plot would be obtained by keeping \mathbf{A}_0 constant and varying m . Assuming that the dashed line with $\lambda = 0.87$ corresponds to $m = 0$, the solid line could be generated by taking $m = 0.103$ and the dotted line by taking $m = 0.207$. In fact, given any fluctuation in this system with $m > 0.149$, the dynamics of the perturbed system will be divergent although the unperturbed system converges. Moreover, even small m can cause a huge increase in convergence time as demonstrated by the dashed and solid curves in the figure.

Generally speaking, even if \mathbf{M} does not take the special form studied above, a majority of positive entries in \mathbf{M} indicate that the average time to solution will very likely be larger than the unperturbed prediction. It is a property of nonnegative matrices that increases to individual elements result in an increase of the largest eigenvalue (Minc, 1998). Even if \mathbf{A}_0 or \mathbf{M} has some negative entries, λ will still be greater than λ_0 on average according to (7) if the average entry of \mathbf{M} is positive.

4.2. Temporal Variability

The shift in largest eigenvalue brought about by \mathbf{M} and described in the previous section can engender delays in the average dynamics, which are still described by either smooth convergence to, or divergence away from, solution. Far more interesting are the deviations from this smooth path brought about by the temporal fluctuations which result in a jagged, unpredictable evolution. As we now show, a high enough level of temporal variation produces a skewed distribution of finishing times and can result in long delays.

Recall that the temporal variation of the fluctuations is represented by the matrix \mathbf{S}_t , which is a matrix of random variables with mean zero. To be as general as possible, we assume that each entry of this matrix is independent of any other. Analogously to the previous section, we also assume for definiteness that the typical size of the fluctuations is proportional to the respective average interaction values. For example, if the element $(A_0 + M)_{35} = 0.1$, then the standard deviation of the (3,5) element of \mathbf{S}_t for any t is equal to $0.1s$, where s is a constant of proportionality. These statements can be concisely expressed by the correlation rule

$$\langle (S_t)_{ij}(S_t)_{i'j'} \rangle = s^2[(A_0 + M)_{ij}]^2 \delta_{i'i'} \delta_{j'j'}, \quad (10)$$

where δ is the Kronecker delta.

Notice that the state of the system is determined by many random matrices multiplied together. This type of stochastic effect is known as multiplicative noise. One aspect of systems with multiplicative noise is that their distribution are approximately log-normal, with parameters proportional to time. For small fluctuations, the variance decreases with t and the log-normal distribution is very similar to a normal distribution, in which events far from the average are very unlikely. However, once the fluctuations exceed a certain threshold, the variance increases with t and the system's log-normal distribution develops a heavy tail at long times, implying that events far from the average are likely.

The log-normal character of the system is shown in figures 4 and 5, below. In figure 4, a histogram showing the work remaining on task 1 in the vehicle interior appearance design project (see Tables 1 and 2, above) is shown at $t = 30$ for various fluctuation strengths. Each plot's distribution is approximately log-normal as shown in the figure by the fitted curves; note the differences in shape as the fluctuation strength increases. In figure 5, the time to solution is plotted for several levels of noise; the stopping criterion is that all components of the state vector be smaller than 0.05. This log-normal distribution of solution times is in agreement with previous results on cooperative problem solving (Clearwater et al., 1991).

In order to understand the conditions under which fluctuations lead to large deviations from the average dynamics, we study the moments of the state vector. In general, the dynamics of systems whose second or third⁴ moments $\langle |\mathbf{x}_t|^2 \rangle$ and $\langle |\mathbf{x}_t|^3 \rangle$ are significantly larger than the powers of the average $|\langle \mathbf{x}_t \rangle|^2$ and $|\langle \mathbf{x}_t \rangle|^3$ are dominated by large deviations from the average.

In our particular case, the moments may become orders of magnitude larger than powers of the average as time goes on. This is because the moments can diverge in time even as the average converges to zero. In this situation the problem solving process is highly susceptible to long delays brought about by large deviations from the average. An example of the deviations brought about by small and large fluctuations is shown in figure 6 below.

We can obtain a mathematical condition on the matrices involved which determines when the deviations described above will occur. To do this, recall expressions (6) for the evolution of the average state, which imply that powers of the average are given by

$$|\langle \mathbf{x}_t \rangle|^p = c^p e^{p \ln \lambda} \quad (11)$$

for large t . Here we are using the notation λ for the largest eigenvalue of $\mathbf{A}_0 + \mathbf{M}$ and we have normalized the right eigenvector \mathbf{u} to have length 1, as is customary. In general, the moment evolution at large t can also be expressed in this form:

$$\langle |\mathbf{x}_t|^p \rangle = c^p e^{t L_p}, \quad (12)$$

where L_p is the p th moment Lyapunov exponent and is always greater than the corresponding exponent in (11). The interesting case described above occurs when $\lambda < 1$, so that the average dynamics are convergent, but $L_p > 1$ for $p = 2$ or 3 so that the second or third moments diverge.

The Lyapunov exponents are generally difficult to calculate. However, (Wilkinson, 2005) gives a measure of the difference between the Lyapunov exponents and the corresponding

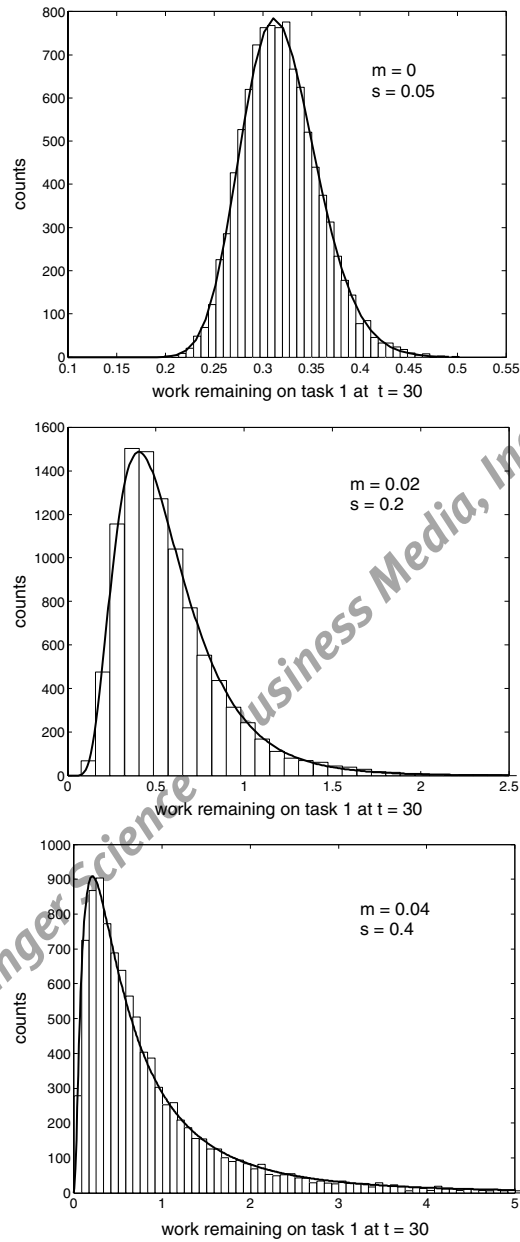


Figure 4. Histograms showing the distribution of the first component of \mathbf{x} at time $t = 30$ over 10,000 simulations of the vehicle interior appearance design project for various fluctuation strengths. The unperturbed WTM for this process is given in Table 2. Recall the notation m for the size of the fluctuation mean and s for the strength of the temporal variation. The superimposed curves are log-normal probability functions calculated using the sample mean and variance and normalized by the bin width and height. The curve parameters are $\hat{\mu} = 0.316, 0.51, 0.603$ and $\hat{\sigma} = 0.123, 0.486, 1.019$ respectively by increasing fluctuation strength.




Figure 5. Distribution of finishing times for various fluctuation strengths over 10,000 simulations of the vehicle interior appearance design project. The curve parameters are $\hat{\mu} = 109.9, 142.61, 191.07$ and $\hat{\sigma} = .0831, .179, .3238$ respectively for increasing fluctuation strength.

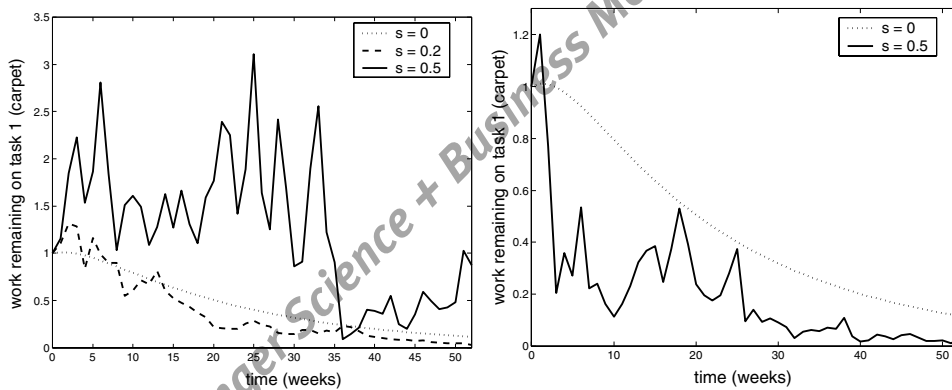


Figure 6. Example of dynamics for various fluctuation strengths. At left, the fluctuation causes a delay in task 1. At right, task 1 proceeds more quickly than an unperturbed treatment would indicate. Recall that the parameter s from (10) measures the strength of the fluctuations.

exponents in the powers of the average:

$$\varepsilon^2 = \frac{s^2}{\lambda^2} \sum_{i,j} v_i^2 [(A_0 + M)_{ij}]^2 u_j^2, \tag{13}$$

where \mathbf{u} and \mathbf{v} are the right and left eigenvectors, respectively, of $\mathbf{A}_0 + \mathbf{M}$ corresponding to λ . Large values of ε^2 indicate that L_p is significantly larger than $p \ln \lambda$. In this case, even a process which converges rapidly on average, i.e. with a low λ , may undergo large deviations and encounter delays. When ε^2 is small, deviations are unlikely unless the convergence of the average is very slow, i.e. λ is close to 1. Since the variance of the (i, j) element of




Figure 7. Stability diagram for a system with $n = 5$. Recall that b is the absolute strength of the fluctuation, while s is the strength proportional to the average size of the entries of $\mathbf{A}_0 + \mathbf{M}$.

the fluctuations is given by $s^2[(A_0 + M)_{ij}]^2$, it is clear from Eq. (13) that the effect of fluctuations is increased as their magnitude increases, as shown in figure 6.

To show analytically how ε^2 affects the dynamics, we consider cases in which the Lyapunov exponent in Eq. (12) can be approximated calculated. When the evolution is discrete and ε^2 small, the moments evolve according to Wilkinson (2005)

$$\begin{aligned} \langle |\mathbf{x}_t|^2 \rangle &= c^2 e^{t(2 \ln \lambda + \varepsilon^2)}, & \text{second moment} \\ \langle |\mathbf{x}_t|^3 \rangle &= c^3 e^{t(3 \ln \lambda + 3\varepsilon^2)}, & \text{third moment.} \end{aligned} \tag{14}$$

These expressions show that the moments will diverge as ε^2 increases.

In the case in which the elements of the fluctuations all have the same variance b^2 , Wilkinson (2005) provides a measure of the critical strength of this variance above which deviations are likely:

$$b_c^2 = \frac{1}{n + \frac{v \varepsilon \lambda^2}{1 - \lambda^2}}, \tag{15}$$

where v is the length of the left eigenvector \mathbf{v} . When the variance of the fluctuations exceeds b_c^2 , the second moment of the system will diverge. By applying the relation (7) we can find a corresponding critical value s_c^2 , and thus determine the stability diagram of the project dynamics. The critical values are plotted for $n = 5$ in figure 7. When $\lambda < 1$ but the fluctuation strength is larger than s_c^2 or b_c^2 , deviations from the average are likely and although the process converges on average, significant delays are likely to occur in particular realizations of the project.

5. Conclusion

In this paper we presented a dynamical model of cooperative problem solving projects consisting of any number of interrelated tasks. The model incorporates the unavoidable fluctu-

ating nature of the interactions in such projects, which reflect the natural variations in the amount of progress made on each task, as well as the varying relevance of information from the partial solutions of each task to any other. While simple, the model captures key aspects of the dynamics, such as the diminishing value of interactions as the tasks near completion, as well as allowing for a satisficing decision of when the project is completed. It also captures the fact that projects with too many tasks do not converge to a desirable solution at all.

By analyzing the stochastic dynamics we showed that as the strength of the fluctuations in task effectiveness increases past a certain threshold, the convergence of a typical process to a solution is very likely to deviate from what an average treatment would indicate. In other words, single problem solving instances are likely to deviate significantly from what one would expect if fluctuations were ignored. In particular, a cooperative problem solving process may be pushed far from its goal, resulting in a significant delay in the time to solution. While this effect is most severe in problems whose average convergence is slow, large fluctuations can also disrupt the smooth convergence towards a solution in any project. This prediction is in agreement with numerous empirical results.

Moreover, over appropriate time intervals the distribution of performance among tasks is log-normally distributed, in agreement with previous empirical work on cooperative problem solving efforts. The appearance of a heavy tailed distribution above a threshold value in task variability implies that repeated observations of the evolution of a project will tend to markedly differ from the average dynamics that other models have analyzed. In other words while there might be instances where time to completion seems to be in agreement with an estimate of the average lifetime of the project, on other occasions project duration will be dramatically longer, with the concomitant aggravation of not being able to predict when such long delays will take place.

These results suggest that in order to have a predictable outcome in large cooperative projects, effort should be made to organize cooperative tasks in a way that minimizes both the instabilities implied by large unstructured groups, and the fluctuating nature of their contributions. This can be achieved by structuring tasks in such a modular or hierarchical way so that they effectively interact with very few others. Furthermore, given the stochastic nature of such processes, one possible mitigating design factor is to keep fluctuations in performance constrained to variations far below their threshold value.

Given the importance of large teams in the design and solution of technical problems and the increasing need for production agility, the study of the dynamics of group problem solving offers new insights into the organization of teams and their efficiency at solving complex tasks.

Appendices

A. *Fluctuations have no Effect on the Average*

As mentioned above, the average dynamics of the system are not affected by the fluctuations. In order to see this, notice that in the discrete case the system evolves according to $\mathbf{x}_t = (\mathbf{A}_0 + \mathbf{M} + \mathbf{B}_t)\mathbf{x}_{t-1}$. Since the components of the matrix \mathbf{B}_t are independent of the state at

previous times,

$$\langle \mathbf{x}_t \rangle = [\mathbf{A}_0 + \mathbf{M} + \langle \mathbf{B}_t \rangle] \langle \mathbf{x}_{t-1} \rangle$$

and we are taking the expected value of \mathbf{B}_t entry by entry, $\langle (B_t)_{ij} \rangle = 0$ for all t, i and j , we have $\langle \mathbf{x}_t \rangle = (\mathbf{A}_0 + \mathbf{M}) \langle \mathbf{x}_{t-1} \rangle$ and, extrapolating back to $t = 0$, we obtain Eq. (5).

B. Dynamics of a Linear System

For any dynamical system specified by the equations

$$\langle \mathbf{x}_{t+1} \rangle = \mathbf{C} \langle \mathbf{x}_t \rangle,$$

where \mathbf{C} is a constant matrix, the state at any time may be expressed in terms of the initial state by writing

$$\langle \mathbf{x}_t \rangle = \mathbf{C}^t \mathbf{x}_0 \quad (16)$$

In the long time limit, the matrix product \mathbf{C}^t is utterly dominated by the contribution from the largest eigenvalue λ of \mathbf{C} for large t . To see this, we diagonalize⁵ \mathbf{C} into the form $\mathbf{C} = \mathbf{P} \mathbf{\Lambda} \mathbf{P}^{-1}$, where $\mathbf{\Lambda}$ is the diagonal matrix of eigenvalues of \mathbf{C} . This simplifies equations (16) because

$$\begin{aligned} \mathbf{C}^t &= \mathbf{P} \mathbf{\Lambda}^t \mathbf{P}^{-1} \\ &= \mathbf{P} \begin{pmatrix} \lambda_1^t & & \\ & \lambda_2^t & \\ & & \ddots \\ & & & \lambda_n^t \end{pmatrix} \mathbf{P}^{-1}, \end{aligned}$$

where the $\{\lambda_i\}$ are the eigenvalues of \mathbf{C} .

In our model, the matrix $\mathbf{C} = \mathbf{A}_0 + \mathbf{M}$ has a dominant largest eigenvalue $\lambda = \lambda_1$. Therefore, when t is large, all terms in the above matrix other than λ^t are completely negligible. The expressions are thus dominated by λ and its eigenvectors \mathbf{u} (corresponding column of \mathbf{P}) and \mathbf{v} (corresponding row of \mathbf{P}^{-1}). Expanding the matrix product with all entries of $\mathbf{\Lambda}$ neglected except the ones with λ , we obtain (6).

C. Log-Normal Distribution

The log-normal probability distribution describes a random variable whose logarithm is normally distributed. The probability density function is

$$P(x) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-\frac{(\ln x - \ln \mu)^2}{2\sigma^2}} \quad (17)$$

where μ and σ are the mean and standard deviation of the normal distribution that describes $\ln x$. The p th moment of the log-normal distribution is given by

$$\langle x^p \rangle = \exp \left(p \ln \mu + \frac{p^2 \sigma^2}{2} \right). \quad (18)$$

In our application, σ^2 and $\ln \mu$ are proportional to t . For large σ^2 the moments diverge for large enough p , causing large deviations from the average.

D. Continuous Time Model

A continuous model is applicable in situations where the rate of exchange of information between people working on different tasks is high compared to the rate at which work is done. This may be the case in a small company, for example, where employees frequently and consistently update each other on the details of their work, for example over lunch or in chance meetings. In this continuum limit one obtains for the governing equation of the model, analogous to Eq. (4), the stochastic differential equation (Øksendal, 2003)

$$d\mathbf{x} = [(\mathbf{A}_0 + \mathbf{M} - \mathbf{I})dt + d\mathbf{S}(t)]\mathbf{x}(t), \quad (19)$$

where $d\mathbf{S} = \mathbf{S}(t)dt$ is a matrix of independent mean-zero Wiener processes with standard deviation proportional to \sqrt{dt} and \mathbf{I} is the identity matrix. For mathematical simplicity, we have assumed here that the noise is white, i.e. that the fluctuations are Markovian.

As in the discrete case, the fluctuations have no effect on the evolution of the average. This can be seen by formally integrating (19):

$$\mathbf{x}(t) = \int_0^t (\mathbf{A}_0 + \mathbf{M} - \mathbf{I})\mathbf{x}(s)ds + \int_0^t d\mathbf{W}\mathbf{x}(s).$$

Since $d\mathbf{W} = \mathbf{B}(t)dt$ and the elements of \mathbf{B} have mean 0,

$$\left\langle \int_0^t d\mathbf{W}\phi(\mathbf{x}, s) \right\rangle = 0$$

for any function $\phi(\mathbf{x}, s)$ and the expected value is given simply by

$$\langle \mathbf{x}(t) \rangle = \int_0^t (\mathbf{A}_0 + \mathbf{M} - \mathbf{I})\langle \mathbf{x}(s) \rangle ds$$

so that

$$d\langle \mathbf{x} \rangle = (\mathbf{A}_0 + \mathbf{M} - \mathbf{I})\langle \mathbf{x}(t) \rangle dt.$$

This equation is the continuous analog of (5).

Since $\mathbf{A}_0 + \mathbf{M} - \mathbf{I}$ is a constant matrix, the average dynamics of this system are given in the long time limit by

$$\langle \mathbf{x}(t) \rangle = e^{(\mathbf{A}_0 + \mathbf{M} - \mathbf{I})t} \mathbf{x}(0)$$

Proceeding analogously to Appendix B, the matrix exponential may be written

$$\begin{aligned} e^{(\mathbf{A}_0 + \mathbf{M} - \mathbf{I})t} &= \mathbf{P} e^{\mathbf{\Lambda} - \mathbf{I}t} \mathbf{P}^{-1} \\ &= \mathbf{P} \begin{pmatrix} e^{(\lambda_1 - 1)t} & & & \\ & e^{(\lambda_2 - 1)t} & & \\ & & \ddots & \\ & & & e^{(\lambda_n - 1)t} \end{pmatrix} \mathbf{P}^{-1}, \end{aligned}$$

where the λ are the eigenvalues of $\mathbf{A}_0 + \mathbf{M}$. In the large time limit, the (1, 1) element of this matrix completely dominates all the other terms since $\lambda = \lambda_1$ is the largest eigenvalue. We thus find

$$\langle \mathbf{x}(t) \rangle = \mathbf{c} \mathbf{u} e^{(\lambda - 1)t}.$$

for the average evolution, which is the continuous analog of Eq. (6). It is apparent from comparing this Eq. to (6) that the behavior of the average dynamics in the continuous case will be qualitatively identical to those of the discrete case.

Acknowledgment

This work was partially supported by the National Science Foundation under Grant No. 9986651.

Notes

1. More precisely, the diagonal elements indicate the percentage of work left incomplete after one time period for each task, as further discussed below.
2. This condition is fulfilled by almost every reasonable interaction matrix. A simple eigenvalue has only one associated eigenvector, and thus its eigenspace is one-dimensional. By dominant we mean that $\lambda > |\lambda_i|$ for all $\lambda_i \neq \lambda$. A dominant eigenvalue is necessarily real, and the dominant eigenvalue of an interaction matrix must be positive. There do of course exist matrices whose largest eigenvalue is not simple, as well as matrices with multiple largest eigenvalues, i.e. $\lambda_1, \lambda_2, \dots, \lambda_n$ such that $|\lambda_1| = |\lambda_2| = \dots = |\lambda_n|$. However, it can be easily shown (Minc, 1998, e.g., or see Wilkinson (2005) for an explicit discussion) that a nonnegative system with several largest eigenvalues splits into independent subsystems each with a simple dominant eigenvalue of equal or smaller magnitude. If some of the elements of $\mathbf{A}_0 + \mathbf{M}$ are negative, then there can be a dominant complex conjugate pair of eigenvalues. In this case the asymptotic behavior is not smooth, but oscillatory. However, the convergence properties are the same with λ replaced by the complex absolute value $|\lambda|$.
3. The case where \mathbf{v} is perpendicular to \mathbf{x}_0 is irrelevant because the values of $\mathbf{A}_0 + \mathbf{M}$ are impossible to exactly determine. In addition, we note that it is possible for elements of \mathbf{v} to be very large; this situation corresponds to matrices whose eigenvalues are highly susceptible to fluctuations. While of great interest, such matrices are

rare and usually not reasonable interaction matrices. This topic is quite complicated and beyond the scope of this paper; see (Golub and van Loan, 1989), or Wilkinson (2005) for a discussion in relation to this particular problem.

4. Large higher-order moments also indicate deviations from the average but with smaller probability.
5. It is not necessary to consider defective (non-diagonalizable) matrices in our model because such matrices form a set of measure zero, and it is impossible to exactly determine the values of the elements of the interaction matrices.

References

- R. Ahmadi, T. Roemer, and R. Wang (2001), "Structuring Product Development Processes," *Eur. J. Oper. Res.*, 130, 539–558.
- Reza H. Ahmadi and Hongbo Wang (1994), *Rationalizing product design development processes*. UCLA Anderson Graduate School of Management Working Paper.
- T.J. Allen (1966), "Studies of the Problem-Solving Process in Engineering," *IEEE Trans. Engrg. Management*, EM-13(2), 72–83.
- Ludwig Arnold, M.M. Doyle, and N. Sri Namachchivaya (1997), "Small Noise Expansion of Moment Lyapunov Exponents for Two-Dimensional Systems," *Dynamics and Stability of Systems*, 12(3), 187–211.
- F. Brooks (1975), *The Mythical Man-month*. Addison-Wesley, Reading, Mass.
- T. Browning, J. Deyst, S.D. Eppinger, and D. Whitney (2002), "Adding Value in Product Development by Creating Information and Reducing Risk," *IEEE Trans Engrg Management*, 49(4), 428–442.
- T.R. Browning and S.D. Eppinger (2002), "Modeling Impacts of Process Architecture on Cost and Schedule Risk in Product Development," *IEEE Transactions on Engineering Management*, 49(4), 428–442.
- Tyson Browning (2001), "Applying the Design Structure Matrix to System Decomposition and Integration Problems: A Review and New Directions," *IEEE Trans. Eng. Management*, 48(3), 292–306.
- M. Carrascosa, S.D. Eppinger, and D.E. Whitney (1998), "Using the Design Structure Matrix to Estimate Product Development Time," In *Proceedings of the ASME Design Engineering Technical Conferences (Design Automation Conference)*, Atlanta, GA.
- K.B. Clark (1989), "Project Scope and Project Performance: The Effect of Parts Strategy and Supplier Involvement on Product Development," *Management Sci.*, 35(10), 1247–1263.
- Kim B. Clark and Takahiro Fujimoto (1991), *Product Development Performance: Strategy, Organization, and Management in the World Auto Industry*. Harvard Business School Press, Boston.
- Scott Clearwater, Bernardo Huberman, and Tad Hogg (1991), "Cooperative Solution of Constraint Satisfaction Problems," *Science*, 254, 1181–1183.
- M. Cusumano and R. Selby (1995), *Microsoft Secrets*. Free Press, New York.
- S.K. Ethiraj and D. Levinthal (2004), "Modularity and Innovation in Complex Systems," *Management Sci.*, 50(2), 159–173.
- D. Ford and J. Sterman (1998), "Dynamic Modeling of Product Development Processes," *System Dynamics Review*, 14(1), 31–68.
- D. Ford and J. Sterman (1999), *Overcoming the 90 percent syndrome: Iteration management in concurrent development projects*. Working paper, Texas A&M University.
- Z. Füredi and J. Komlós (1981), "The Eigenvalues of Random Symmetric Matrices," *Combinatorica*, 1(3), 233–241.
- J.R. Galbraith (1977), *Organizational Design*. Addison-Wesley, Reading, Mass.
- Gene H. Golub and Charles F. van Loan (1989), *Matrix Computations*. The Johns Hopkins University Press, Baltimore, 2nd edn.
- A. Griffin (1997), "The Effect of Project and Process Characteristics on Product Development Cycle Time," *J. Marketing Res.*, 10(1), 24–35.
- Albert Y. Ha and Evan L. Porteus (1995), "Optimal Timing of Reviews in Concurrent Design for Manufacturability," *Management Sci.*, 41(9), 1431–1447.
- M. Hammer (1996), *Beyond Reengineering*. Harper Business, New York.

- Tad Hogg and Bernardo A. Huberman (1993), "Better than the Best: The Power of Cooperation," In Lynn Nadel and Daniel Stein, (eds.), *1992 Lectures in Complex Systems*, volume V of *SFI Studies in the Sciences of Complexity*, pp. 165–184. Addison-Wesley, Reading, MA.
- Tad Hogg, Bernardo A. Huberman, and Jacqueline M. McGlade (1989), "The Stability of Ecosystems," *Proceedings of the Royal Society of London*, B237, 43–51.
- Bernardo A. Huberman and Natalie S. Glance (1993), "Evolutionary Games and Computer Simulations," *PNAS*, 90, 7716–7718.
- Bernardo A. Huberman and Tad Hogg (1995), "Communities of Practice: Performance and Evolution," *Computational and Mathematical Organization Theory*, 1(1), 73–92.
- William C. Ibbs (1997), "Quantitative Impacts of Project Change: Size Issues," *Journal of Construction Engineering and Management*, 123(3), 308–311.
- Yan Jin and Raymond Levitt (1996), "The Virtual Design Team: A Computational Model of Project Organizations," *Computational and Mathematical Organization Theory*, 2(3):171–195.
- N.R. Joglekar (2001), *Data collected at Factory Mutual Insurance Company*, Norwood, MA.
- F. Juhasz (1982), "On the Asymptotic Behavior of the Spectra of Nonsymmetric Random (0,1) Matrices," *Discrete Mathematics*, 41, 161–165.
- Stephen J. Kline (1985), "Innovation is not a Linear Process," *Res. Management*, 28(4), 36–45.
- Viswanathan Krishnan, Steven D. Eppinger, and Daniel E. Whitney (1997), "A Model-Based Framework to Overlap Product Development," *Management Sci.*, 43(4), 427–438.
- Christoph Loch, C. Terwiesch, and S. Thomke (2001), "Parallel and Sequential Testing of Design Alternatives," *Management Sci.*, 47(5), 663–678.
- C. Mar (1999), *Process improvement applied to product development*. MS thesis, MIT.
- C.D. McDaniel (1996), *A linear systems framework for analyzing the automotive appearance design process*. MS thesis, MIT.
- Jürgen Mihm, Christoph Loch, Bernardo A. Huberman, and Dennis M. Wilkinson. "Hierarchies and Problem Solving Oscillations in Complex Organizations," preprint.
- Jürgen Mihm, Christoph Loch, and Arnd Huchzermeier (2003), "Problem-Solving Oscillations in Complex Projects," *Management Science*, 49(6), 733–750.
- Henryk Minc (1988), *Nonnegative Matrices*. John Wiley and Sons, New York.
- P.W.G. Morris (1982), "Project Organizations: Structures for Managing Change," In A.J. Kelley, ed., *New Dimensions of Project Management*, pp. 155–174. Lexington, MA. Heath and Co.
- P.W.G. Morris and G.H. Hugh (1987), *The Anatomy of Major Projects*. Wiley, Chichester.
- Bernt K. Øksendal (2003), *Stochastic Differential Equations: An Introduction with Applications*. Springer, Berlin, 6th edn.
- F. Peña-Mora and M. Li (2002), "Dynamic Planning and Control Methodology for Design/Build Fast-Track Construction Projects," *Journal of Construction Engineering and Management*, 127(1), 1–17.
- J.S. Reel (1999), "Critical Success Factors in Software Projects," *IEEE Software*, 16(3), 18–23.
- N. Repenning, P. Gocalves, and L. Black (2001), "Past the Tipping Point: The Persistence of Firefighting in Product Development," *Calif. Management Review*, 43(4), 44–63.
- J.W. Rivken and N. Siggelkow (2003), "Balancing Search and Stability: Interdependencies Among Elements of Organizational Design," *Management Sci.*, 49(3), 290–311.
- T.A. Roemer, R.H. Ahmadi, and R.H. Wang (2000), "Time-Cost Trade-Offs in Overlapped Product Development," *Oper. Res.*, 48(6), 858–865.
- Robert P. Smith and Steven D. Eppinger (1997), "Identifying Controlling Features of Engineering Design Iteration," *Management Sci.*, 43(3), 276–293.
- M.E. Sosa, S.D. Eppinger, and C.M. Rowles (2003), "Identifying Modular and Integrative Systems and their Impact on Design Team Interactions," *Trans. ASME*, 125, 240–252.
- Gilbert Strang (1976), *Linear Algebra and Its Applications*. Academic Press, New York.
- Tetsuo Tamai and Akito Itou (1997), "Requirements and Design Change in Large-Scale Software Development: Analysis from the Viewpoint of Process Backtracking," in *Proceedings of the 15th international conference on Software Engineering*, pp. 167–176, Los Alamitos, California. IEEE Computer Society Press.
- C. Terwiesch and C.H. Loch (1999), "Managing the Process of Engineering Change Orders," *J. Product Innovation Management*, 16(2), 160–172.

- C. Terwiesch, C.H. Loch, and A. de Meyer (2002), "Exchanging Preliminary Information in Concurrent Engineering: Alternative Coordination Strategies," *Organization Science*, 13(4), 402–419.
- Stefan Thomke (1997), "The Role of Flexibility in the Development of New Products," *Res. Policy*, 26, 105–119.
- J. Thompson (1978), *Organizations in Action*. McGraw-Hill, New York.
- K. Ulrich (1995), "The Role of Project Architecture in the Manufacturing Firm," *Research Policy*, 24(3), 419–440.
- T. van Zandt (1999), "Decentralized Information Processing in the Theory of Organizations," In M. Sertel, ed., *Contemporary Economic Issues*, pp. 125–160, London, MacMillan.
- Daniel E. Whitney (1990), "Designing the Design Process," *Res. Engineering*, 2:3–13.
- Dennis M. Wilkinson (2005), "Moment Instabilities in Multidimensional Systems with Noise," *Eur. J. Phys. B*, 43:221–242.
- Ali Yassine, Nitin Joglekar, Dan Braha, Steven Eppinger, and Daniel Whitney (2003), "Information Hiding in Product Development: The Design Churn Effect," *Res. Engineering Design*, 14:145–161.

Bernardo A. Huberman is a Senior HP Fellow and Director of the Information Dynamics Laboratory. He is also a Consulting Professor of Physics at Stanford University. For the past ten years he has concentrated on understanding distributed processes and on the design of mechanisms for information aggregation and the protection of privacy as well as market-based distributed resource allocation systems.

Dennis Wilkinson is a recent graduate of Stanford University with a doctorate in Physics, and has accepted a position in the Department of Defense. His research interests include dynamics of social networks and other stochastic systems, information extraction from large, complex networks, and techniques in distributed computing.