PPST Report

Ben Tsou

September 11, 2006

The control of quantum phenomena is being actively pursued today for two general reasons: (1) producing products that are inaccessible to more conventional means; (2) revealing the fundamental nature of interactions in atoms and molecules. A quantum system without control typically evolves in an undesirable fashion. Therefore, a tailored control field is introduced to redirect the quantum dynamics towards a desirable, perhaps even an optimal outcome. Quantum optimal control theory (OCT) and optimal control experiment (OCE) methods have been successfully employed to identify optimal controls for a variety of quantum systems. To seek the maximum degree of control, a quantum control objective must be chosen. Regardless of the objective function chosen, the dynamics of the system is always guided by the unitary evolution operator U, which is a functional of the control field $C(t)$.

Two commonly used objectives are the optimization of state-to-state transitions and the generation of arbitrary unitary transformations. For the former, the goal is to maximize the probability $P_{if} = |U_{if}|^2$ of making a transition from some initial state $|i>$ to some final state $|f>$. For the latter, the goal is to generate (by controlling $C(t)$) a unitary operator U that best matches a given unitary operator W. This can also be stated as minimizing the cost function $J = ||U - W||$, the Hilbert-Schmidt distance between U and W.

Good solutions for quantum optimal control problems have been found for both objectives. However, OCT studies show that the search effort required for the the $||U - W||$ objective greatly exceeds that for the P_{if} objective. In particular, the $||U - W||$ objective becomes exponentially more difficult to search with rising dimension while the P_{if} objective is largely unaffected by rising dimension. A useful tool in analyzing these problems is the concept of a quantum control landscape. The landscape critical topology of these two objectives is well understood. However, the difference in searchability between them is still not well understood.

A grid-walking algorithm called the 'LGW Method' has been developed to analyze the structure of these landscapes. For each run of the 'LGW Method', a grid-size δ is chosen. This grid-size defines a cartesian grid on which a 'walker' can move; the 'walker' can move by a step size δ in each of the 2N coordinate directions (where N is the dimension of the control space). At each iteration of the algorithm, the 'walker' moves in the most favorable direction (towards the objective value) until it gets stuck.

The 'LGW Method' gives insight into the texture of the landscapes in the neighborhood of critical point regions. Our findings show that the texture of the P_{if} landscape does not change with rising dimension. On the other hand, the texture of the unitary landscape grows tremendously more complex with rising dimension. Thus, the difference in searchability between the two objectives can be at least partially explained by the differences in texture elucidated by the 'LGW Method'.

The 'LGW Method' offers two measures of the texture of landscapes: convergence probability and average terminal value (ATV). To calculate the convergence probability for a given δ size, one first sets a threshold value. Then many sample runs are conducted and the percentage of runs that exceeded the threshold value is the convergence probability. The average terminal value (for a given δ size) is simply the average of the values at which the runs terminated (the 'walker' got stuck).

It was found that regardless of the threshold value chosen, the critical δ size for approximately half the runs to converge (convergence probability of 50 percent) remains constant with increasing dimension for the P_{if} objective. On the other hand, the critical δ size decreases exponentially with increasing dimension for the unitary objective. This means that for higher dimensions, a finer mesh or gridsize is required to reach the threshold objective value.

It was also found that as a function of δ, the approximation ATV $\sim c(N)\delta^2$ holds for both objectives. Here, the coefficient c is a function of the dimension N. For the P_{if} objective, c is nearly constant with dimension. However, for the unitary objective, c increases super-exponentially with dimension. This means that for higher dimensions, a finer mesh or gridsize is required to reach a given average value.

Both of these analyses are strong indicators that the texture of the landscape for the unitary objective is much more complex than that for the P_{if} objective. The texture of the unitary landscape depends heavily on the dimension whereas the texture of the P_{if} landscape is largely invariant to dimension.

Although we know that the texture of these landscapes influences the search effort for finding optimal controls, the precise relationship is unknown.

A fundamental investigation of the 'LGW Method' itself will give us more insight into this relationship. To analyze the algorithm, it helps to run it on 'toy' functions where the analysis is much simpler. For several of the simplest 'toy' functions, an analytical formula has been derived for both the convergence probability and ATV as a function of δ. Interestingly, certain features of the convergence probability and ATV plots do not seem to depend on what function is used, and are therefore intrinsic to the algorithm. For example, the plot of the convergence probability vs. δ follows an S-shaped curve for almost all functions tested. Also, ATV $\sim c(N)\delta^2$ holds not only for the two quantum objectives, but for many of the 'toy' functions as well. Understanding these patterns will enhance the LGW Method's use in elucidating the structure of quantum control landscapes. In fact, there is no reason the 'LGW Method' must be limited to quantum landscapes. The 'LGW Method' can be run on any landscape and may be useful when applied to similar landscape problems that arise in other areas.