Chapter 7

Logical Complexity

At the heart of the definition of complexity is the difficulty of making generalizations regarding the relationship between any two variables out of the context of all related variables. However, algorithmic complexity states that only purely chaotic systems are incompressible. If the elevator system is not purely chaotic, the relationship between complexity measures and performance is theoretically expressible in some compressed format, *i.e* some generalizations are possible.

Rather than perform a complete analysis for every combination of variables, a detailed analysis for both potential and actual complexity measures is performed on simple systems - ones close to the current naval weapons elevator system in operation. From that point, we build on the analysis as required and discuss the salient differences between a configuration and previously analyzed configurations. This method is applied to all complexity measures, starting with the logical complexity.

The modeled system that most closely represents the existing weapons elevator systems is a system with conventional shafts (bi-directional with a single dedicated carriage) with serial carriage and hatch/door operations and the use of interlocks. While not critical, it is beneficial to define a system size and connectivity equivalent to the existing system in order to establish a baseline configuration for comparison of possible improvements and to serve as model validation. This validation is performed in detail in Appendix B. However, establishing the size and connectivity of the existing system is difficult for several reasons. Foremost is the fact that the actual weapons handling system is controlled by an ordnance handling officer who relies on subjective judgment rather than adhering to a predefined rule set to direct operations. While the physical arrangement and connectivity of the actual and model systems may be identical, the *effective* connectivity may be different because of the handling officer's control. For instance, the handling officer may delay the loading of a carriage from one queue in order to continue loading of items of a type found only in another queue so these items are eliminated from consideration.

Additional reasons for the difficulty in defining a specific system size and connectivity are the ambiguous definition of queues and the allocation of resources. Many operations are required to support a carrier air wing and, since space is limited on a carrier, queues are established where room permits. The number, locations, sizes, and connectivities of queues may all be different for any given strike-down scenario. Additionally, the distribution of resources in a queue may vary - where they are shared in one situation, they may be dedicated in another.

Given these differences, it is arguable that the existing system consists of several independent queue-shaft-magazine circuits (1-1-x size systems), where x is the number of magazines served by the shaft. In its most complicated operation, the existing system is a 1-2-x system, where some fraction of x magazines are shared by both shafts. The 1-1-1 system is that described by the Navy as the baseline scenario for defining target operation cycle times and will be discussed specifically when examining performance/behavior relationships.

7.1 Simple Conventional Systems

As stated in Chapter 6, logical complexity is defined as the ratio of the length of the logical evolution to the temporal evolution length and is intended to identify the level of logical 'activity' in an evolution. Figure 7.1 shows the relationship between the logical complexity and throughput for all conventional SIL systems. Throughput is in units of items per minute, where one item per carriage is carried per trip. This measure of throughput is used throughout this work. The general outline is a skewed triangle with left, right, and upper boundaries. At first glance, the correlation between complexity and throughput is apparently meaningless. For a large range of logical complexities, there also exists a large range of throughputs. However, the general trend resulting from the common skew of the left and right boundaries is that more logically active evolutions have increased performance.

The correlation between complexity and throughput for all evolutions is 0.921, indicating a relatively strong relationship. The value of the correlation must be made in light of the data represented however. We will see later that the limits on maximum complexity and throughput are dependent on system size - specifically that smaller systems have lower maximum complexity and throughput. The set of all evolutions considered therefore affects the correlation. If more smaller, simple systems are considered, the relationship between complexity and performance will be different than if only larger systems are evolved. The set of evolutions considered is limited practically by the speed and memory of the computers used for simulations, but the limitations of computers should not have any influence on the realities of the relationships between variables. Because the collective use of all evolutions is invalid with respect to drawing conclusions regarding the relationships between complexity and performance, analyses are performed on multiple evolution sets, where each set is comprised of configurations of identical size (*e.g.* 1-1-1 systems, 2-2-2 systems, 3-4-2 systems, etc…).

Our analysis begins with the simplest of systems, the single queue, shaft, and magazine system. The throughput for this system is predictable, being the inverse of the sum of all individual carriage cycle times. Similarly, the logical complexity is also predictable as the ratio of the number of unique states in a normal cycle (4) to the time required to complete a cycle. Since there is only one 1-1-1 valid configuration, there is only a single data point in Figure 7.2. While this plot is not particularly useful in establishing a correlation between behavior and performance, the case provides an easily analyzable evolution and baseline absolute values. Both the values of logical complexity and throughput, when compared to Figure 7.1 are seen to be minimal. Intuitively, the minimal complexity matches our intuition. The description of the system made above is very short and the system's evolution is quite predictable, with no interdependencies. As a result, we should expect complexity to be minimal.

Since there is only a single carriage, both the compressed and full state trajectories and system state histories are simple to follow[1]. The compressed state trajectory shown in Figure 7.3 shows the system cycling through four states. These are the four possible states which occur for any given carriage, excluding the case of traveling to an unspecified queue, which only occurs in this case at the last logical step.

The compressed and full state histories are presented in Figure 7.4 and identify the same pattern - a simple repetitive cycle through four states: loading at the queue, traveling down to the magazine, unloading in the magazine, and traveling empty to the queue.

Larger systems with a single carriage all behave like a 1-1-1 system and theoretically have the same throughput and logical complexity. The state trajectories and state histories may appear slightly more complex in these cases however. Figure 7.5 shows the evolution of a 1-1-4 system with a 20-20-20-40 queue distribution. The different cycles are a result of the carriage visiting different magazine numbers, but the evolution is characteristically simple and has the same throughput as the evolution with a single magazine. The logical complexity is also identical as it does not identify unique patterns in an evolution, only the state changes.

In larger size systems with greater than one carriage, variations in connectivity are possible and a relationship between logical complexity and throughput begins to develop. Figure 7.6 shows the correlation between the logical complexity and throughput for all 2-2-2 systems (of which there are 66 evolutions)[2].

The trend line of the right bound seen previously in the set of all results in Figure 7.1 is partially evident in these results, as is the divergence from the minimal values.

Although there are 66 valid evolutions of 2-2-2 size systems, less than 66 data points are visible in Figure 7.6. The discrepancy exists for two reasons. The first is due to the resolution of the plot. Since the points used in the plot have a real size, nearby points share the same space and overlap, obscuring the true distribution. The second reason is that many

[1]The carriage state history is identical to the log_2 of the system state history because there is only a single carriage

[2]The scale for the complexity is kept constant for all system sizes to facilitate comparisons

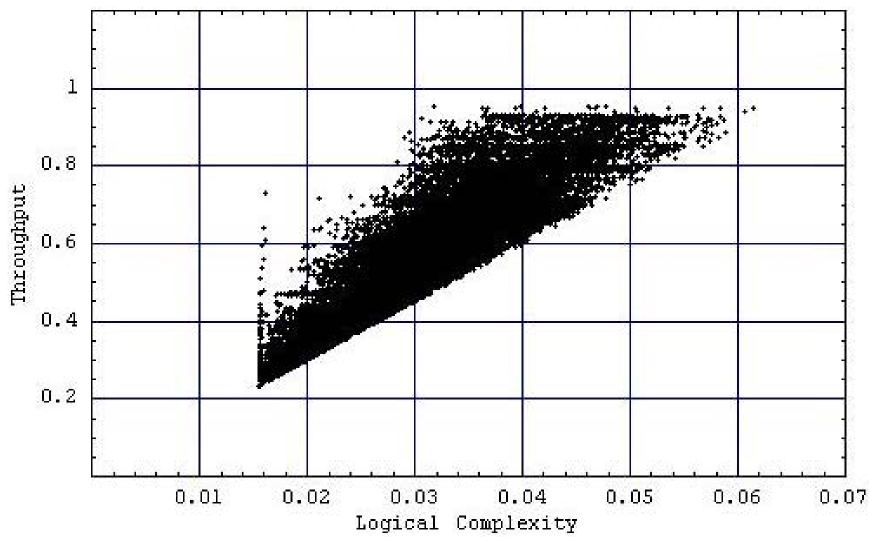

Figure 7.1: The distribution of all evolutions for conventional SIL systems with respect to logical complexity and throughput. 546919 data points are represented. The variables have a correlation of 0.921.

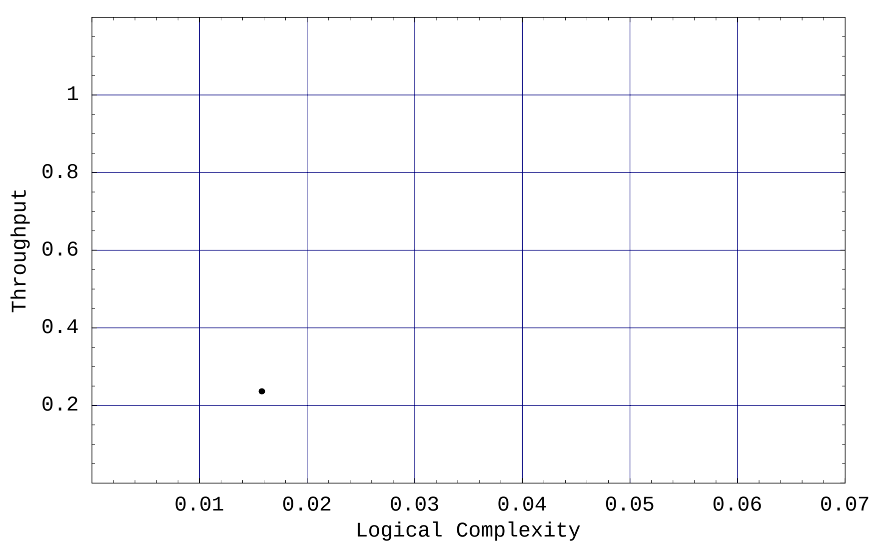

Figure 7.2: There is only a single 1-1-1 configuration and evolution. When transients are ignored, the logical complexity and throughput are both minimal values for all evolutions.

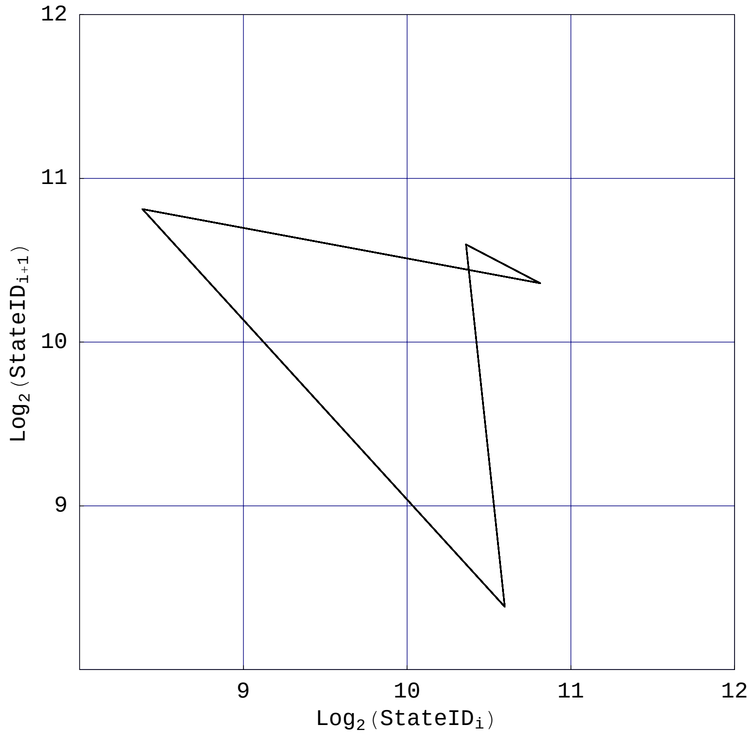

Figure 7.3: The compressed state trajectory for configuration 3 1-1-1 shows a simple cycle through 4 states.

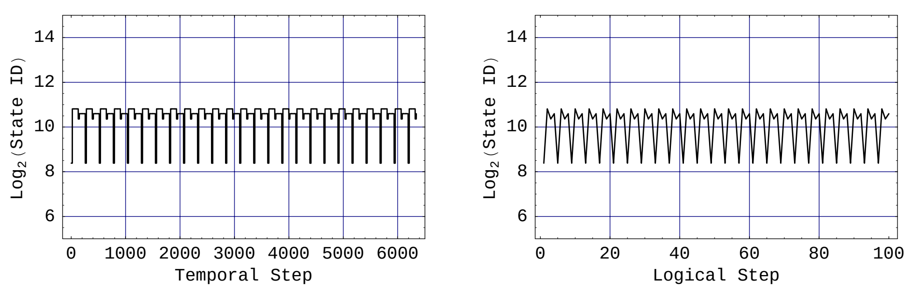

(a) (b)

Figure 7.4: The full and compressed state histories for the 3 1-1-1 evolution. These histories also identify a simple cycle through 4 unique states.

points have identical throughput and complexity and differentiation is impossible without identifying the frequency of values. The three dimensional frequency landscape, discussed in Section 5.6, accounts for these visualization problems and represents a truer picture of the behavior - performance relationship. The frequency landscape for 2-2-2 size systems is shown in Figure 7.7 and indicates that many of the evolutions share identical, or nearly identical,

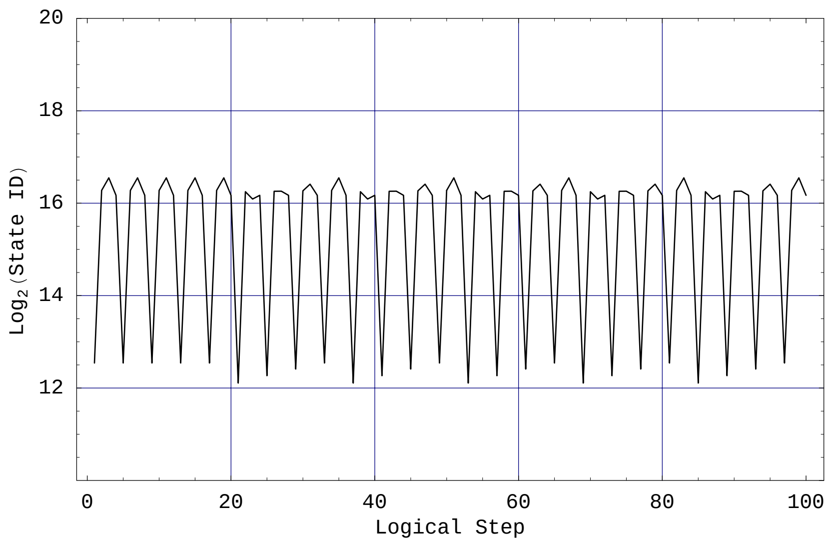

Figure 7.5: The compressed state history of the single 1-1-4 system with a 20-20-20-40 queue distribution. The evolution appears more complex than that for the 1-1-1 configuration, but the behavior is simple in both cases and throughput and logical complexity are identical.

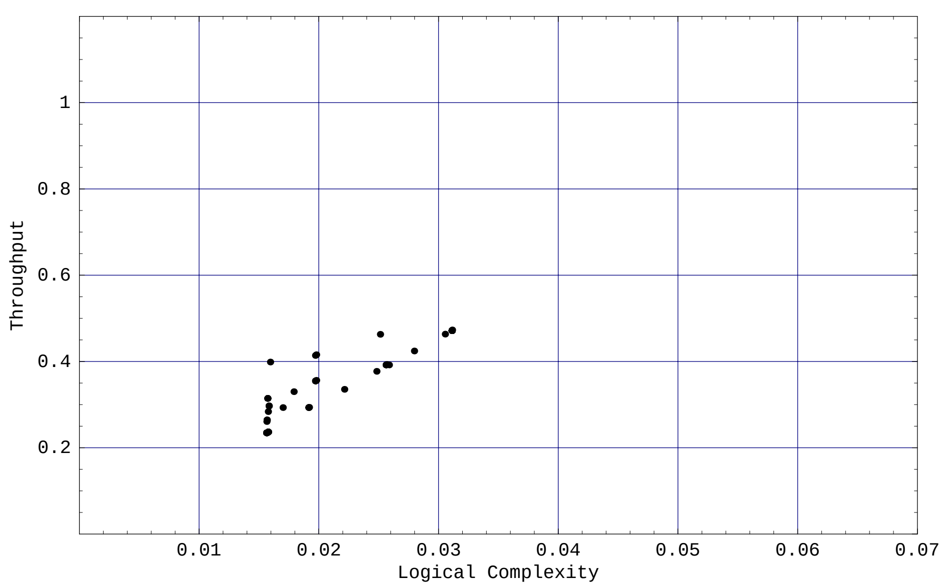

Figure 7.6: The relationship between logical complexity and throughput for 2-2-2 size systems. With more queues, shafts, and magazines, variations of connectivities are possible and result in a distribution of evolutions.

values[3].

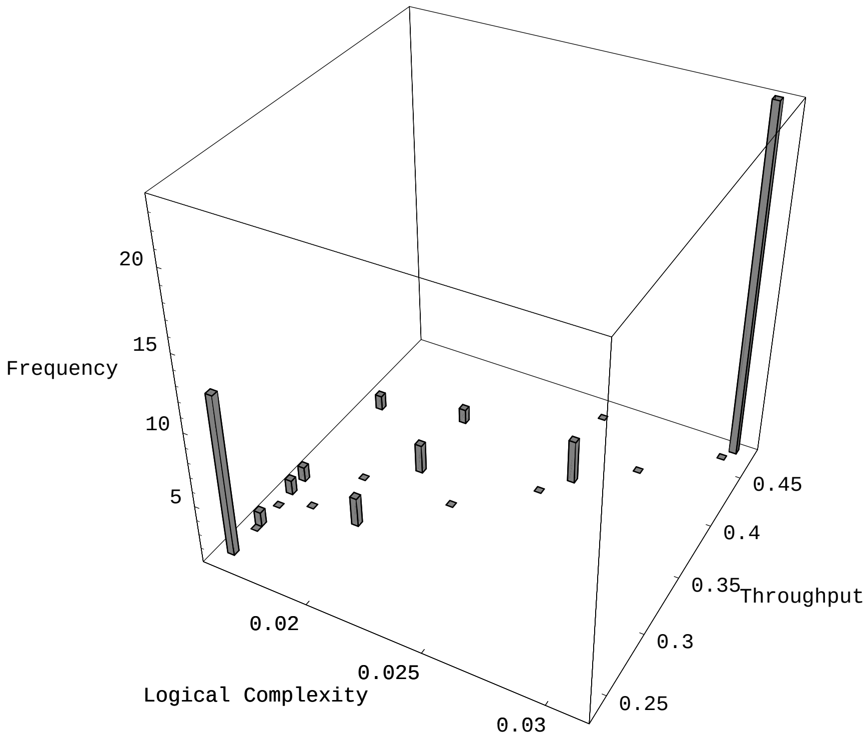

Figure 7.7: The frequency landscape for all 2-2-2 evolutions. The various frequencies indicate the presence of close or identical complexity and throughput values.

The frequency landscape shows that results are clustered at the intersection of the maximal values of logical complexity and throughput and the same minimal values of logical complexity and throughput resulting from the simple 1-1-1 evolution. Collectively, these results comprise 53% of all results (35% for maximal values and 18% for minimal values). This distribution of results helps explain the correlation of 0.931 and the nature of the least squares fit of $0.0731 + 12.86 LogicalComplexity$ which passes near both peak values.

The frequency landscape has the additional purpose of identifying the evolutions of most interest for detailed analysis. Concentrating on the most common evolutions yields the greatest characterization of the complexity/throughput relationship. Two of the 2-2-2 configurations

[3]The axes of the frequency landscapes are scaled according to the minimum and maximum values of complexity and throughput to offer the most detail, unlike the two dimensional relational plots, which use fixed scales

that represent the set of evolutions with minimal logical complexity and throughput are 118 (100-0) and 247 (100-0). These configurations are selected because they are the configurations with the least and most physical connectivity and therefore indicate a bound on the relationship between connectivity and behavior. For both configurations, 100% of the items in both queues are bound for the first magazine. The physical connectivity of configuration 118 is presented in Figure 7.8.

$$(SQ) = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$$

Figure 7.8: The physical connectivity for configuration 118 2-2-2. Since only the second shaft is connected to the first magazine, the first shaft is essentially non-existent for a queue distribution of (100-0).

The SM incidence matrix indicates that only the second shaft can access the first magazine. Since both queues only contain items destined for magazine 1, shaft 2 is effectively non-existent and all items from both queues are transported by the first shaft, making this scenario essentially equivalent to a 2-1-1 size system. The operation of this evolution is evident from the compressed state history in Figure 7.9, showing the same fundamental simple cycles as those for the 1-1-1 system in Figure 7.4(b), where the first set of cycles represents transport from the first queue and the second set corresponds to deliveries from the second queue.

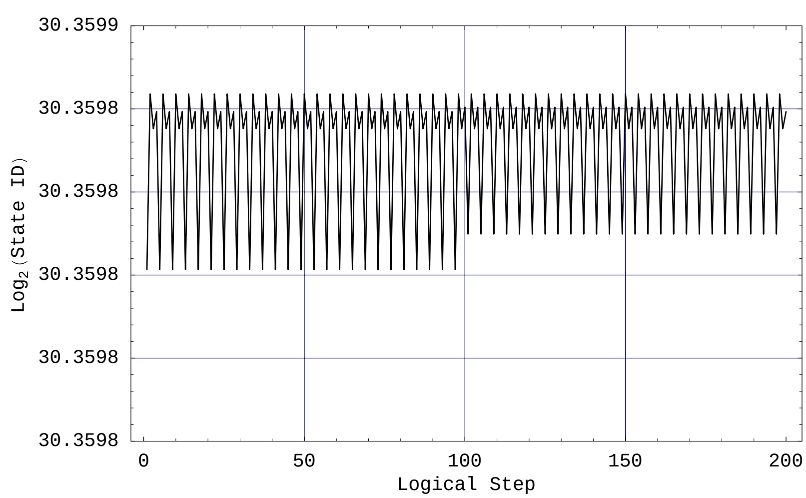

Figure 7.9: The compressed state history for the 118 2-2-2 (100-0) evolution with minimal logical complexity and throughput. The pattern is identical to that for the 3 1-1-1 evolution, only the second carriage transports items from both queues, yielding two sets of simple patterns. The range of state identification values is low because the bits used to describe the unused first carriage are more significant in the state code.

The evolution is clearer in the carriage evolution histories, shown in Figure 7.10, because corresponding bits for each carriage have the same significance. The carriage histories show the two sets of cycles corresponding to travel to the two queues by the first carriage as well as a flat line corresponding to the static second carriage.

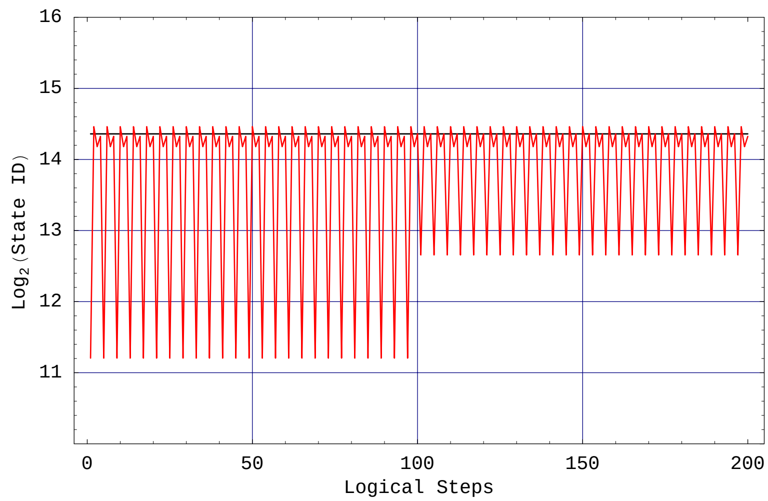

Figure 7.10: The compressed evolution histories for individual carriages for the 118, 2-2-2 configuration with a (100-0) queue distribution. The static first carriage exists in one state and is represented by a horizontal line while the utilized second carriage runs through two simple patterns.

Configuration 247 2-2-2 with a (100-0) queue distribution behaves in exactly the same manner as 118 2-2-2 and 31 1-1-1, again with the second carriage performing all deliveries while the first carriage remains static. It is interesting to note that configuration 247 is more completely connected than 118, with incidence matrices shown in Figure 7.11. Despite this increased physical connectivity however, 247 performs identically to 118 *for this queue distribution.* The identical results of 247 and 118 for the (100-0) queue distribution does not imply that 247 and 118 behave identically in all situations. This is far from the case. Configuration 247 delivers all items for all possible queue distributions and has the highest throughput and logical complexity for several queue distributions. In contrast, 118 generally has low logical complexity and completely transports all items in only a single evolution - the evolution corresponding to the (100-0) queue distribution. It must be recalled that the goal here is not necessarily to compare configurations, although it is important to show the range of complexity and performance possible for various configurations. The primary goal at this point is to illustrate how configurations behave and what sort of evolutions result in certain regions of complexity and throughput.

Configuration 103 2-2-2 (0-100) represents the simplest system from the largest set of evolutions sharing common values of throughput and logical connectivity in Figure 7.7. The physical connectivity of 103 2-2-2 is shown in Figure 7.12.

Configuration 103 delivers all items despite the sparse connectivity for the (0-100) queue

$$(SQ) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 1 & 2 \\ 1 & 2 \end{pmatrix}$$

Figure 7.11: Configuration 247 2-2-2 has greater connectivity than 118 2-2-2, but the values of their logical complexities and throughputs are identical for (100-0) queue distributions.

$$(SQ) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

Figure 7.12: The incidence matrices for 103 2-2-2. Despite sparse physical connectivity, 103 has maximal logical complexity and throughput.

distribution because both shafts have access to the second magazine, which is the destination of all items. Since each queue uses a unique shaft to reach the second magazine, the shafts are essentially independent except when sharing magazine resources. Figure 7.13 shows the compressed state history for this evolution. The presence of a single repetitive cycle is evidence of the carriage independence as the carriages evolve to simple phase-locked cycles with a phase lag. The phase lag occurs in the first evolution cycle when both carriages vie for access to magazine 2 to unload. Since the magazine resources can not handle both carriages, a time delay occurs as the second carriage waits for access (carriage 1 gains access as a tiebreaker). This time delay effectively creates two independent 1-1-1 circuits with identical cycle times, and is evident in Figure 7.14 which shows the carriage evolution histories. Each carriage has its own simple cycle, but both share the same period, collectively creating the single repetitive *system* cycle evident in Figure 7.13. With both carriages in operation, twice as many items are delivered than in the 1-1-1 system in the same time, resulting in a throughput approximately twice that of the simple 1-1-1 configuration (0.471 vs. 0.237). The fact that the measured throughput of 103 2-2-2 is less than twice the throughput of 31 1-1-1 is attributed to the delay which occurs as one carriage waits for access to the magazine and introduces the phase lag. If the two carriages operated completely independently with different queues and magazines, or if enough resources were present to prevent carriage interaction, the throughput would be exactly double, although fully independent circuits would result in a non-robust configuration where queues would be restricted to containing only items for the magazines they were connected to.

Because of the phase lag between individual carriage cycles introduced by the sharing of resources in the evolution of 103 2-2-2, the average number of state changes per system cycle increases to approximately twice that for the evolutions of 31 1-1-1, 118 2-2-2, and 247 2-2-2. The logical complexity therefore suggests the evolution is twice as complex. The evolution is still a simple repetition however, and it would appear as if logical complexity fails in this case as a measure of behavior as the complexities are very similar - "this global cycle composed of these states occurs this many times" describes both evolutions. However, in the light of algorithmic complexity, the logical complexity has meaning. Despite the simple repetitive system pattern of 103 2-2-2, additional information is required to describe the evolution,

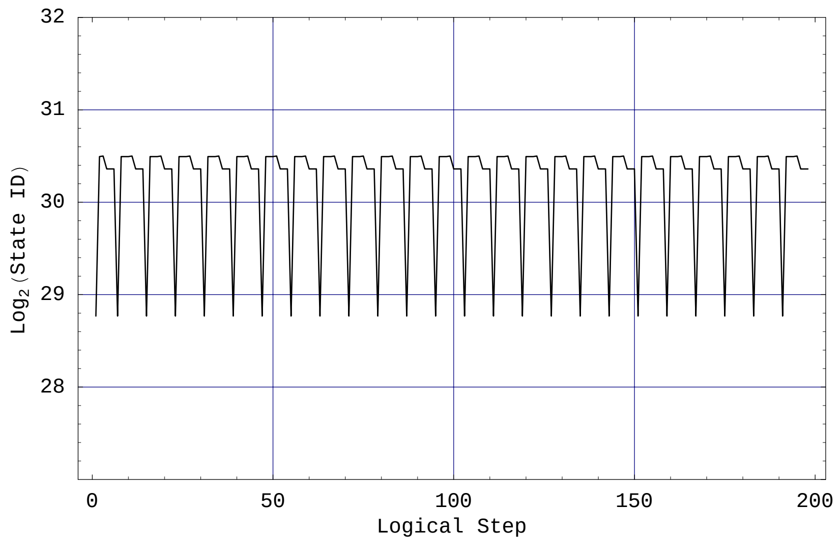

Figure 7.13: The compressed state history of 103 2-2-2 (0-100) shows a single repetitive pattern for the system states. Although both carriages are in operation, they are phase locked with a phase lag, resulting in a single system pattern that contains twice the system states per system cycle as a 1-1-1 system.

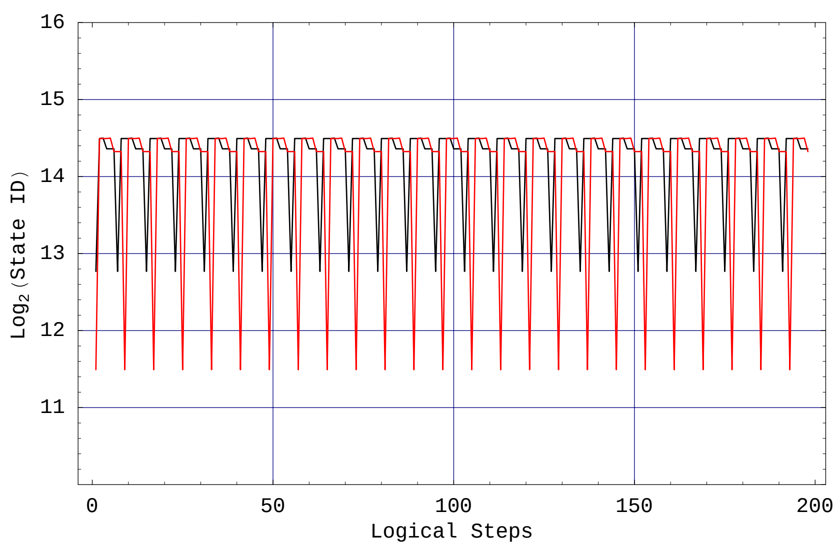

Figure 7.14: The carriage state histories for 103 2-2-2 (0-100) show the constant phase lag between individual carriage cycles. After the initial time delay in the first cycle, the carriages are effectively independent and no longer interact.

most notably the description of the phase lag that occurs at the beginning of the evolution. Additionally, a description of the sequencing of state changes is required for a complete description of the evolution. For each state change, it is necessary to declare that carriage 1 is in state x when carriage 2 is in state y. For a complete evolution description, doubling of the complexity may be a valid estimation and, when transients are included, may even be considered an underestimate.

All evolutions of configuration 255 2-2-2 share the same logical complexity and throughput as 103 2-2-2 (0-100). 255 is a unique configuration however, as it is the only 2-2-2 configuration with complete physical connectivity. Regardless of the item distributions in the queues, the completely connected configuration achieves the greatest throughput and logical connectivity of evolutions of this size. Despite its identical performance and complexity, the evolutions of the completely connected configuration may be distinct depending on the queue distribution. As an example, consider the evolutions of 255 with (0-100) and (60-40) queue distributions.

The carriage state histories for the (0-100) and (60-40) queue distributions are presented in Figure 7.15 and illustrate the differences in behavior that are possible with the same physical architecture but different queue distributions. The evolution of a (0-100) queue distribution shows two sets of simple phase locked cycles separated by a phase lag. Initially, the first carriage removes an item destined for magazine 1 from queue 1. Carriage 2 also selects an item for magazine 1, but from the second queue. The initial selections creates the same phase lag seen in the 103 2-2-2 evolution as carriage 2 must wait for access to magazine 1 as the first carriage unloads. The phase lag introduced at the start of the evolution means the carriages are effectively independent as in the 103 evolution. However, unlike 103, the carriages visit the same magazines *and* queues. The evolution therefore consists of two stages as both carriages empty the first queue, then remove items from the second queue.

With the (60-40) queue distribution, the evolution begins identically to that for the (0-100) queue distribution. However, because of logic specifying that a carriage must select the most common item type in any queue, carriage 2 selects an item destined for magazine 2 in its fourth cycle after carriage 1 removed a magazine 1 type item, leaving more type 2 than type 1 items. The items are subsequently removed from the first queue at identical rates until the queue is emptied. The same process occurs in the second half of the evolution as the carriages remove items from the second queue.

The effect of the phase lag and carriage independence on logical complexity and throughput is the same as that for 103. The time delay introduced by the initial wait for magazine access means the carriages are effectively unaffected by the presence of each other. Therefore, regardless of the differences in evolution histories, the systems visit the same number of logical states per system cycle. With no time delays beyond the initial delay, the evolution lengths are identical and therefore the logical complexities are identical.

Based on one interpretation of algorithmic complexity, the complexity of the two 255 evolutions should be equivalent. Since the number of state transitions is the same, as is their sequencing, the amount of information required to describe the evolutions should also be the

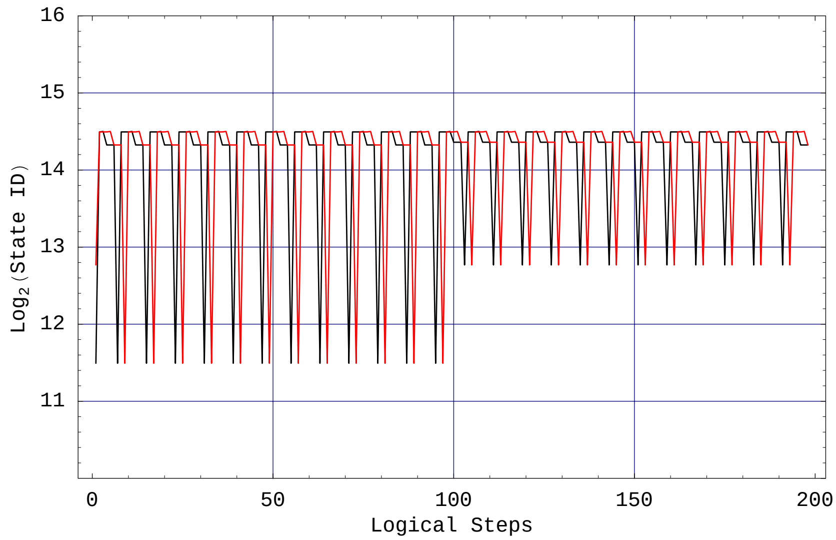

(a)

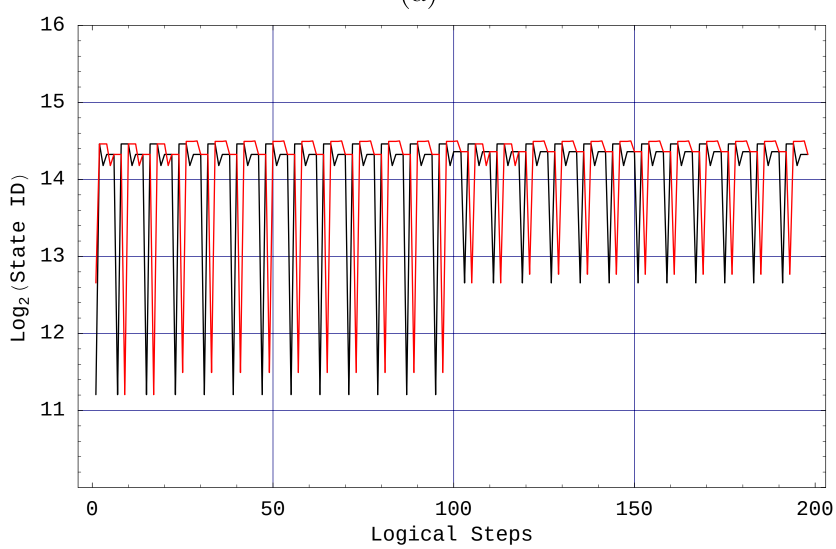

(b)

Figure 7.15: The carriage state histories of the completely connected 2-2-2 configuration are distinct for different queue distributions. In this case, logical complexity and throughput are identical for the evolutions with a (a) (0-100) queue distribution and (b) (60-40) queue distribution.

same. The fact that different states are used should have no bearing on the information content of the description as long as the description is serial - the system visits this state, then this state, then.... However, the evolution of 255 with a (60-40) queue distribution results in a larger set of states occurring throughout the evolution and in a maximally compressed evolution description, the evolution with a (60-40) queue distribution may therefore appear to require more information and have a greater complexity.

The set of 2-2-2 configurations analyzed characterize the extremes of the right bound evident in Figure 7.6 and represent the majority of results for 2-2-2 size systems. At these extremes, evolutions are fundamentally identical. At the lower extreme, only a single carriage is in operation because of limits of the physical connectivity, the queue distribution used, or a combination of both. At the upper extreme, carriages are phase locked with a lag and operate independently. Between these extremes, behavior is a mixture of the two. Configuration 119 with a (40-60) queue distribution provides an example scenario.

$$(SQ) = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$$

Figure 7.16: The incidence matrices for 119 2-2-2. With a queue distribution of (40-60), the evolution is a mixture of behaviors.

The incidence matrices for this configuration are presented in Figure 7.16 and indicate a fairly well connected system. However, the first shaft only connects the second queue to the second magazine. The second shaft, which accesses all queues and magazines, therefore has a greater load, depending on the queue distributions. The disproportionate access is evident in the carriage state histories in Figure 7.17. The evolution begins in a manner similar to evolutions corresponding to maximal throughput and logical complexity, with both carriages transporting items to the same magazine (in this case, magazine 2) and introducing a phase lag between carriage cycles. Since the first carriage can only carry items to the second magazine, it continues to pull from queue 2 until the queue is empty. After these items are transported, the first carriage stops. The second carriage takes items from the first queue, alternating between items bound for magazines 1 and 2 to keep the ratio of item types constant. When the first queue is empty, the second carriage transports the remaining items in the second queue to the second magazine, which the first carriage could not transport. The result of the first shaft's limited connectivity is a shutdown approximately half way through the evolution (approximately 2/3 through the logical evolution), which decreases throughput. Now the evolution behaves like those corresponding to minimal throughput and complexity. Since only one carriage is in operation in a simple repetitive cycle in the second half of the evolution, the number of logical steps per system cycle is half that of when both carriages are in operation with a phase lag, reducing the logical complexity of the evolution. All of the evolutions in Figure 7.6 on the right bound between the extremes share this characteristic. Because of a combination of limited connectivity and the queue distribution, one carriage shuts down during the evolution, reducing throughput and logical connectivity.

The near vertical line defining the left bound of Figure 7.6 represents a substantial fraction of the total number of 2-2-2 evolutions, which is visible in the frequency landscape of Figure 7.7. The lower vertex of this line corresponds to the case of a single operational carriage exemplified by the evolution of 247 with a (100-0) queue distribution. An example of an evolution at the other vertex is configuration 102 2-2-2 (40-60). The physical connectivity of 102 is presented in Figure 7.18.

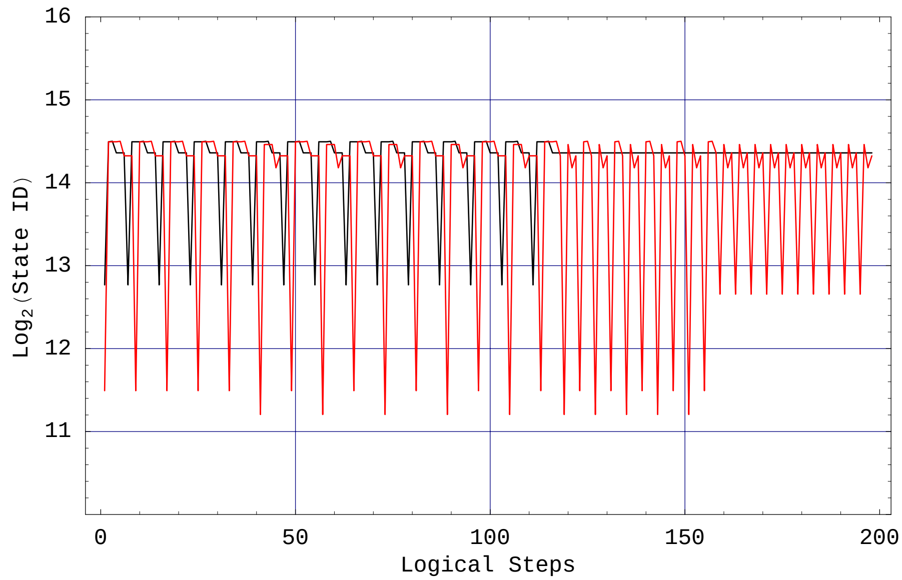

Figure 7.17: The carriage state histories for 119 2-2-2 (40-60) present a mixture of behaviors corresponding to minimal and maximal throughput and logical complexity. Carriage 1 has limited connectivity and shuts down mid way through the evolution, turning a complex evolution into a simpler one.

$$(SQ) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

Figure 7.18: The minimal incidence matrices in a 2-2-2 size system occur in configuration 102. As a result, the carriages act like two independent 1-1-1 circuits.

This configuration is unique because it is the least possibly connected valid configuration and represents a set of two fully independent queue-shaft-magazine circuits. The first shaft can *only* see the second queue and magazine and the second shaft can *only* see the first queue and magazine. With no possible carriage interdependencies, the configuration is therefore equivalent to two independent 1-1-1 configurations acting synchronously.

The carriage state histories for this evolution appear in Figure 7.19. The two carriages operate in simple cycles that are in-phase because they are independent and no time delay ever occurs as a result of resource sharing. Because the carriage cycles are in phase, the system goes through the same number of logical steps per system cycle as a single carriage operating alone. The logical complexity is therefore identical to a single carriage system. Throughput is approximately doubled, because twice as many carriages are in operation. However, throughput falls short of twice the single carriage system throughput because one of the carriages halts after it empties its queue of all items it can transport. Because of its simple connectivity, the carriages carry only the fraction of items in their queue specified by the queue distribution. In the case of the (40-60) queue distribution, the first carriage carries 60% of the items in the second queue and the second carriage transports 40% of the items in the first queue, leaving half of the total number of items in the system undelivered. With identical cycle times, the second carriage therefore stops operation at 2/3 of the total evolution time. From this result, it is apparent that configuration 102 is only capable of transporting half of the items in the system. Extrapolating to larger systems, it is evident that for the configuration with the lowest possible connectivity, only $\frac{1}{n}$ of the total number of items will be delivered, where n is the number of magazines in the configuration. The 102 (40-60) evolution also reveals that, while the logical complexity of all 102 evolutions remains constant regardless of the queue distribution, the throughput is variable. For queues with an even distribution of item types, $\frac{1}{n}$, the throughput is maximal as all carriages halt at the same time. The lowest throughput occurs when only one item type is present in all queues which effectively shuts down all carriages but one. The distribution of evolutions along the left boundary of Figure 7.6 with common logical complexity but various throughput is therefore seen to be comprised of these sparsely connected configurations and their mimics, with carriage termination times and therefore throughput dependent on the ratio of item types in the queues.

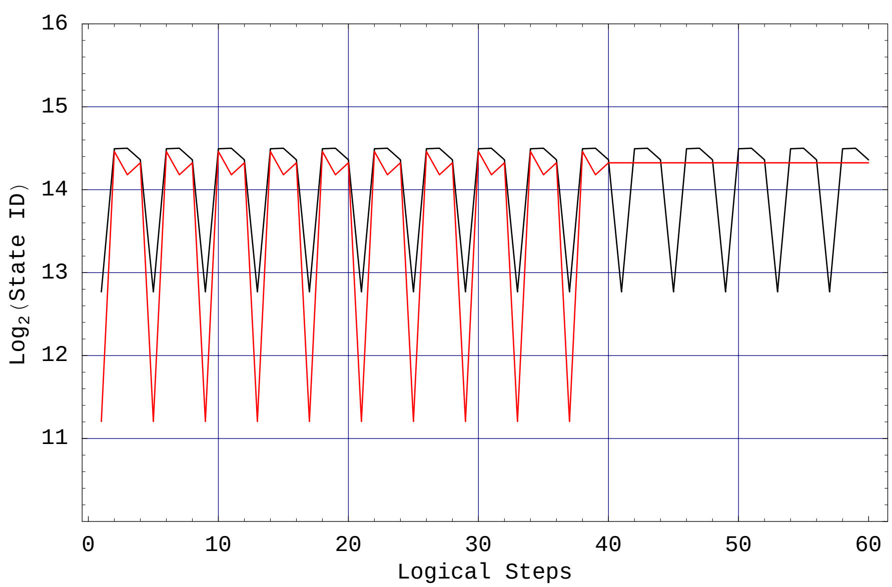

Figure 7.19: The carriage state histories for 102 2-2-2 (40-60) reveal that, for a sparsely connected configuration, logical complexity is a constant, but throughput varies depending on the queue distribution.

The logical complexity of configuration 102 is nearly identical to the logical complexity of a configuration with a single operational carriage exemplified by configuration 247 (100-0). Intuitively, a similar value of complexity appears to be partially correct for these evolutions. Even though there are two carriages in operation in 102 (40-60), they operate independently and in phase. Additional information is required to describe the trajectory of each carriage. However, no additional information is required to describe the sequencing of states like in the description of the two phase lagged carriages in configuration 255. A direct comparison of these evolutions may be invalid however, when system robustness is taken into account.

7.2 Evolution Sets

The sets of evolutions presented thus far have consisted completely of all non-halting evolutions. Halting evolutions are defined as those in which the system logic freezes a carriage in place, leaving deliverable items undelivered. This condition is distinct from an evolution with incomplete item transfer, where undelivered items are undeliverable because of absent connections between queues and item destinations.

While halting evolutions are valid and represent realistic scenarios, they are removed from analysis because the values of throughput and complexity are based on the evolution to the point of termination. The values do not reveal the true nature of a halting evolution and characterization of the performance-complexity space is misleading - throughput in a halting evolution is arguably zero, with deliverable items left undelivered. For similar reasons, it is arguable whether to separate evolutions with incomplete item transfer from the set of all non-halting evolutions. With respect to robustness, surely a configuration that delivers all items is more robust than a configuration that leaves some fraction of the same set of items undelivered. Yet the throughput and logical complexity (and other complexity measures) of these configurations can be identical - the incomplete evolution length is shorter in proportion to the number of undelivered items. As with the results of halting evolutions, the throughput and complexity of incomplete evolutions are misleading, failing to reveal the complete nature of the evolution and to properly characterize the performance-complexity relationship. At the same time, these evolutions represent valid scenarios. If queues are completely controllable, incompleteness is a non-issue as the items which are undeliverable from certain queues would never be included in those queues. Following this argument, the undeliverable items in the incomplete evolutions can be considered as virtual items that are not actually present except in configurations where they are deliverable. Removal of incomplete evolutions therefore results in incomplete characterization of the full range of *possible* evolution behaviors, even if they are somewhat misleading. To account for the opposing arguments for inclusion of incomplete evolutions, analyses of the complete set of non-halting evolutions and the set of evolutions with complete item transfer ("complete evolutions") are performed, with the recognition that incomplete evolutions do not fully reflect all evolution information and that the set of complete evolutions represents a more robust set when queues are uncontrollable.

In an optimal search, the non-halting and complete evolution sets described are larger than they should be. If an evolution of a physical configuration with one queue distribution fails (with respect to halting or incompleteness, depending on which set in which we are interested), then all evolutions of that configuration should be removed from the set (for incomplete evolutions, we are again assuming that the queue is uncontrollable). The remaining configurations are not always necessarily always robust however. Our measure of robustness is based on the set of queue distributions considered in the evolutions. Not halting or complete delivery for these queue distributions is necessary, but not sufficient for inclusion of a configuration into the respective sets. For many configurations, it is undecidable whether any configuration will ever result in a halting or incomplete evolution for any set of queue distributions without explicit evolution of the configuration for all of those queue distributions. Since the complete set of queue distributions has been shown to be quite large, a complete demonstration of robustness is practically impossible. The set of queue distributions used is therefore only intended to represent a significant variety and serve as a go-nogo test for robustness.

The decision to include or exclude configurations for which at least one evolution fails with respect to halting or complete delivery is based primarily on what we are attempting to characterize. If the goal is to characterize the range of possible behaviors of evolutions that meet a certain criterion, then all evolutions that meet the criterion are included in the set, regardless of the other evolutions of the same configuration. If our goal is characterization of possible behaviors in the search for optimal configurations or identification of traits of optimal configurations, then only robust configurations are included.

While in general we are more interested in characterization of the full range of behaviors, characterization of robust sets is important, refining the set and indicating the requirements for consistent inclusion in the set. They also help establish a correlation between performance and complexity - generalizations are possible from the requirements for robustness and the aggregate performance and complexity.

A comparison of the values and distribution of throughput and complexity of sets of evolutions of varying degrees of robustness with respect to completeness with all non-halting and all complete evolutions in 2-2-2 configurations yields some initial correlations. For 2-2-2 configurations, there are six unique queue distributions. The level of robustness of a configuration is determined by counting the number of evolutions of that configuration that have complete item transfer. Of the 66 2-2-2 evolutions, 40 are complete. Of that number, 4 have only one complete evolution and 6 configurations (36 evolutions/6 possible queues) have complete evolutions for all queue distributions. The distribution of evolutions in the complexity-throughput space for these configurations is presented in Figure 7.20.

For configurations with a single complete evolution, the evolutions lie at the minimal logical complexity/minimal throughput and maximal logical complexity/maximal throughput values. The evolutions of robust configurations with complete item transfer for all queue distributions also exist at these extremes, but also run nearly linearly along the entire length of the right logical complexity/throughput bound, suggesting a stronger correlation and higher aggregate complexity. The low complexity/high throughput evolutions (evolutions off the right boundary line) found in the set of all non-halting evolutions are absent from the set of the most robust configurations. As a result, the correlation of logical complexity and throughput for the most robust configurations of 0.997 is greater than the correlation for all non-halting evolutions of 0.931. Additionally, the mean logical complexity for the most

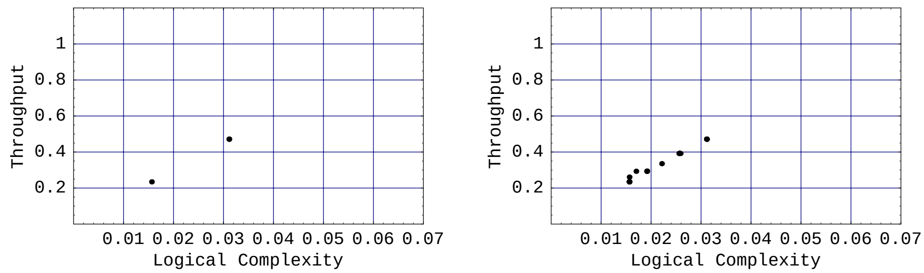

(a) (b)

Figure 7.20: The distributions of 2-2-2 evolutions with (a) 1 complete evolution and (b) 6 complete evolutions. Configurations with 1 complete evolution exist at the minimal logical complexity/minimal throughput and maximal logical complexity/maximal throughput points while the most robust configurations lie along the entire lower boundary.

robust configurations is greater than the mean for all evolutions (0.0267 versus 0.0235) as is the mean throughput (0.407 versus 0.375). While it first may appear that the absence of low complexity/high throughput evolutions in the set of the most robust configurations should result in a lower mean throughput, many of the low complexity/low throughput values are also absent, which is not evident in a comparison of two-dimensional relationships.

Based on these results, it may initially appear that we can make a general statement regarding the correlation of throughput and logical complexity: since the most robust and *adaptable* configurations have a higher mean complexity, then adaptability is directly related to complexity. Similarly, the most adaptable configurations are also the best performers. We have seen this result before - genetic algorithms evolving cellular automata to rule sets resulting in complex and adaptive dynamics capable of computation for instance. Of course, any general statement at this point is tentative for several reasons. First, these results are drawn from the 2-2-2 configuration set, which is a relatively small set with only 66 configurations. The sizes of the sets of complete evolutions and most robust evolutions differ by only 4 evolutions, the number of configurations with only one complete delivery. Second, and more importantly, the correlation and mean values depend on the distribution of evolution sets, which may vary relatively for different evolution size sets. The differences between the distributions of different sets of evolutions is not completely apparent in a comparison of two dimensional plots because of the resolution of the plots and the presence of repeated values. To account for these limitations, three-dimensional frequency landscapes and histograms of cross sections of complexity and throughput are used for comparisons along with two-dimensional distributions.

Figure 7.21 shows the two-dimensional distributions of evolutions for three 2-2-2 size system evolution sets: the set of all non-halting evolutions, the set of complete evolutions, and the evolutions of the most robust configurations. Figure 7.22 presents the three-dimensional frequency landscapes for the same sets. The complexity and throughput histograms for these sets are shown in Figure 7.23.

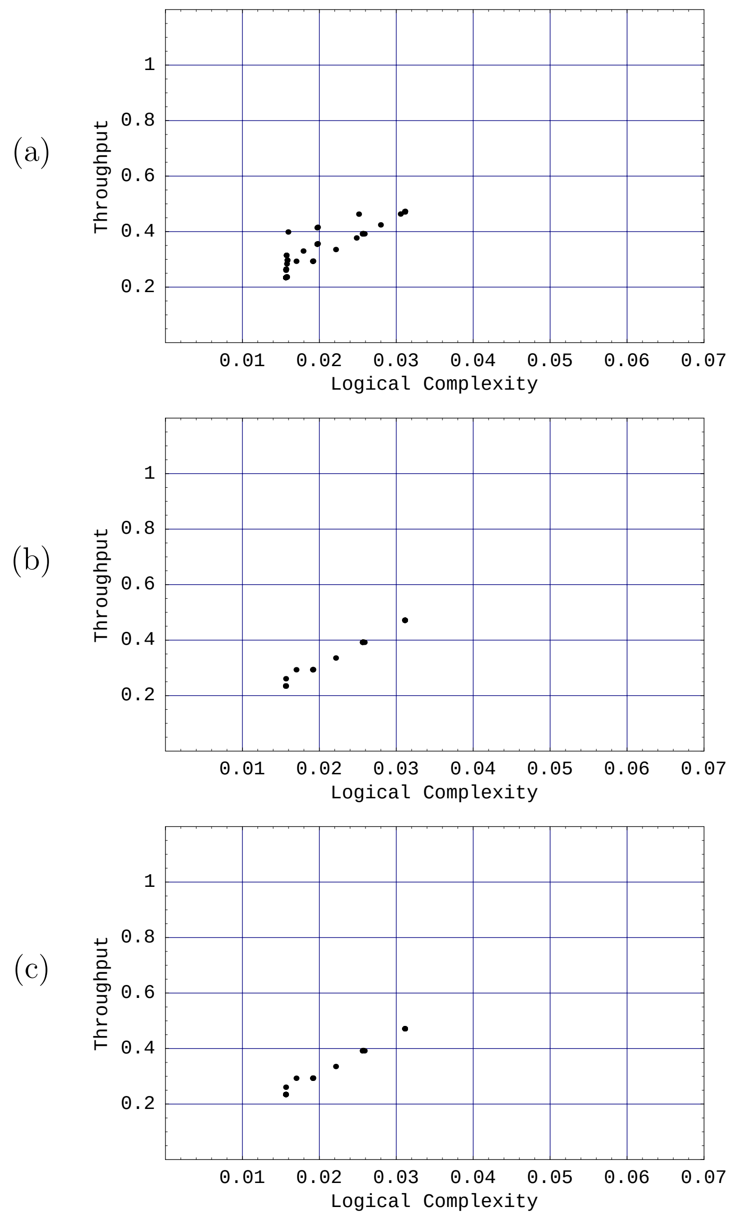

Figure 7.21: The two-dimensional distribution of evolutions for (a) non-halting evolutions, (b) complete evolutions, and (c) the evolutions of the most robust 2-2-2 configurations. The differences between the distributions of evolutions of different sets are not completely apparent.

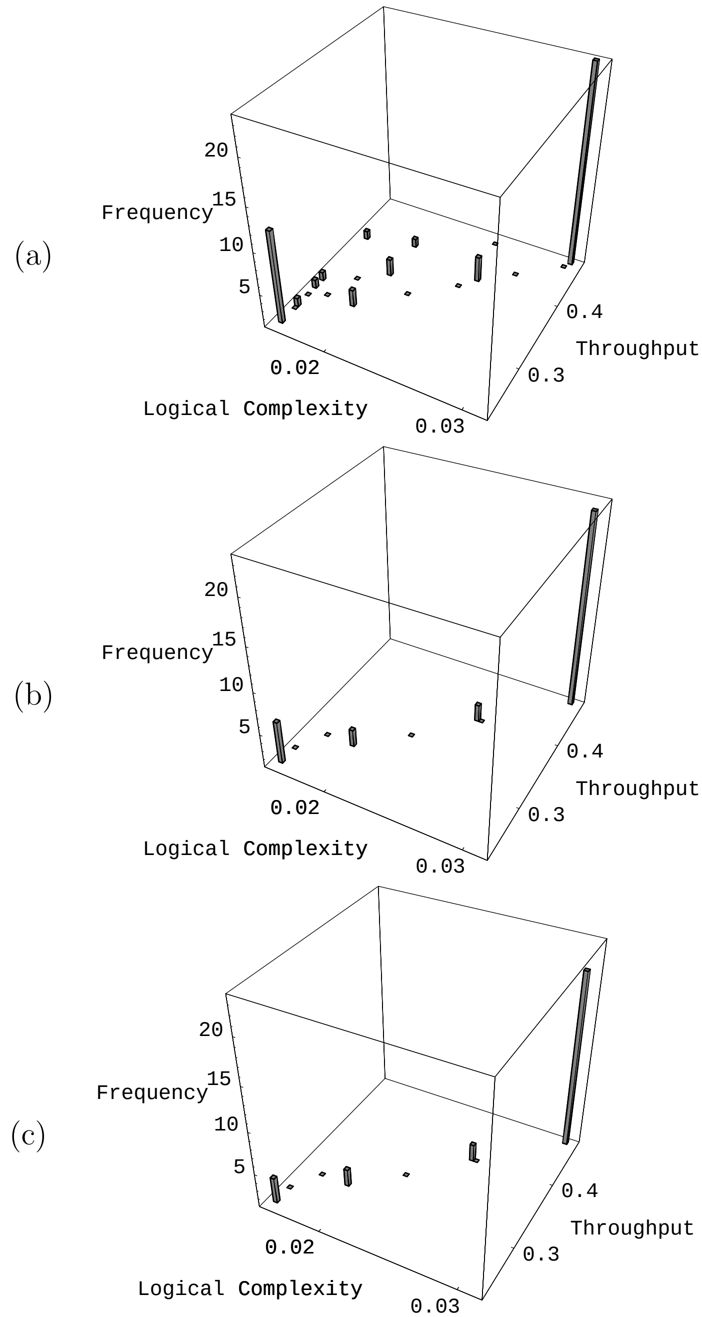

Figure 7.22: The three-dimensional distribution of evolutions for (a) non-halting evolutions, (b) complete evolutions, and (c) the evolutions of the most robust 2-2-2 configurations. The shape of the distribution is less evident than for corresponding two-dimensional plot, but the differences between evolution sets is apparent. The frequency scales are based on the maximum frequency of all three sets to emphasize any differences.

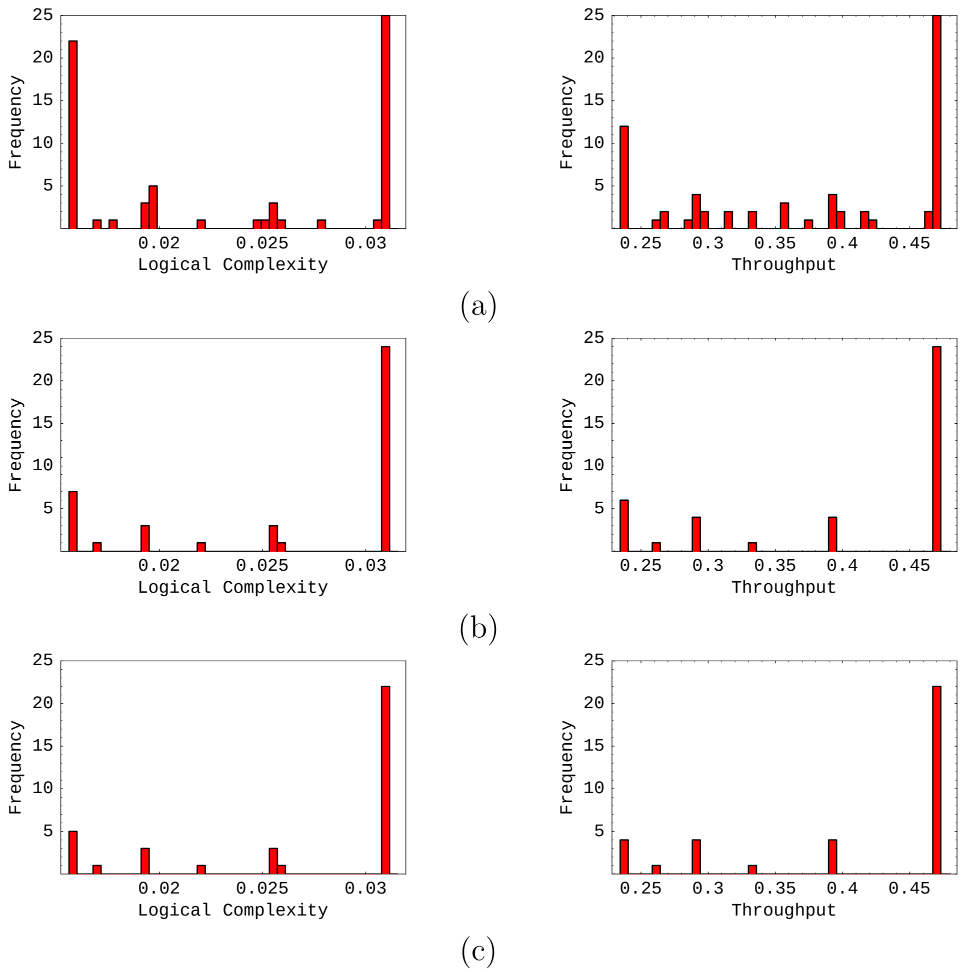

Figure 7.23: The histograms of the evolutions of cross sections of logical complexity and throughput for (a) non-halting evolutions, (b) complete evolutions, and (c) the evolutions of the most robust 2-2-2 configurations. The histograms provide similar information to the three-dimensional frequency landscapes, but the relative frequencies are more apparent.

A comparison of the frequency landscapes and histograms of the complete evolutions and the evolutions of the most robust configurations with the set of non-halting evolutions further illustrates the differences in the distributions that leads to higher aggregate throughput and complexity for the most robust configurations (the set of complete evolutions and evolutions of the most robust configurations differ by only four evolutions). Although the absolute number of high complexity evolutions is lower for the most robust configurations, the ratio of the highest to lowest complexity evolutions increases substantially when the logically simpler evolutions are removed. The differences in the sets of evolutions have a similar, but lessened effect on throughput. Fewer minimal throughput evolutions are removed with the remainder corresponding to moderate throughput values. These moderate throughput values correspond to the low complexity evolutions like 102 2-2-2 (40-60) where a 2 shaft system acts like two independent 1 carriage systems and are guaranteed to have incomplete delivery.

One of the most apparent features of the histograms is that none of them are normal. However, the use of the mean implies that we are dealing with normally distributed data. Therefore, characterization of the distributions using their mean values has a lower significance. The mean does characterize the differences between various subsets of evolutions, like the complete and robust evolution sets. Based on the unique characteristics of sets, it is possible to make general statements regarding complexity and throughput with the mean.

7.3 Mimicry

In the description of the evolutions of configurations 118 2-2-2 and 247 2-2-2 with (100-0) queue distributions represented in Figure 7.10, we saw that only one carriage was in operation during the evolution. For 118, the first carriage carried all items while for 247, only the second carriage was in operation. For this specific queue distribution, the evolutions of these configurations appear identical to the single 1-1-1 configuration, 31, and the similarity is reflected in their values of logical complexity and throughput, which are identical to the simplest system. Evolutions like these, which are equivalent to evolutions of smaller systems and do not utilize all shafts and magazines are referred to as mimics.

Mimicry results from a combination of the physical connectivity and the queue distribution. We can think of these two forms as mimicry with respect to carriages and mimicry with respect to magazines. There is no mimicry with respect to queues because all queues must contain some set of items and be connected to at least one shaft in a valid configuration, so all queues specified in a configuration description are always used.

A requirement for mimicry is a queue distribution with at least one item type absent. If the physical connectivity of the configuration is such that all carriages are utilized to some extent and simply do not visit the magazines corresponding to the absent item types, mimicry is defined to be with respect to magazines. The system acts as if there are as many magazines as there are present item types, but all carriages specified are utilized. This condition occurs multiple times for each configuration. For 2-2-2 configurations with two magazines, it occurs twice for each configuration ((0-100) and (100-0) queue distributions).

In an evolution that exhibits mimicry with respect to carriages, there is at least one carriage that is connected only to those magazines for which there are no items in the queues. In these configurations, the combination of the physical connectivity and queue distribution keep the carriage(s) idle and the configuration mimics a configuration with as many magazines as non-zero item types and as many shafts as non-zero shaft utilizations. Mimics with respect to carriages result in both evolutions with complete item transfer, like 118 and 247, and incomplete item transfer. The set of all non-halting evolutions and the set of complete evolutions can therefore both contain mimics.

In the analyses of 2-2-2 evolutions, we saw that configuration 102, the configuration with the lowest physical connectivity, is essentially two independent 31 1-1-1 circuits. It may initially appear that this configuration is a mimic even with the inclusion of all item types since it is fundamentally no different than the simpler system it emulates. By the definition of mimicry however, the 102 configuration with an all non-zero queue distribution is not a mimic and represents a unique 2-2-2 size configuration. The attribute of the 102 configuration that makes it uniquely 2-2-2 and not a mimic is that the throughput would not be double that of the simple system if not for both carriages. So, while it may represent the most trivial 2-2-2 configuration, 102 is still distinctly 2-2-2.

If the queues are viewed as uncontrollable, then mimicry with respect to magazines becomes less important and evolutions with absent item types and static magazines represent valid scenarios that should be included in the characterization of unique evolutions for a given system size. Mimicry with respect to carriages however, is a function of both the queue distribution and the physical connectivity and is therefore more closely related to the configuration definition. The set of configurations or evolutions that are unique to a given system size are therefore defined as those that do not exhibit mimicry with respect to carriages, but may exhibit mimicry with respect to magazines. The subset of evolutions or configurations that are unique to a system size excludes evolutions or configurations that mimic smaller systems and therefore illustrates a clearer picture of the range of behaviors possible for that system size. The attributes of unique system sets illustrate the effects in changes in system size on complexity and performance. Additionally, from the unique sets and comparisons of subsets of the unique sets, correlations between complexity and performance are established that are, in a sense, truer than the correlations made from sets that include mimics because of the elimination of redundancies.

For the same reasons of establishing correlations and misleading values of complexity and throughput that resulted in the creation of the non-halting evolution and complete evolution subsets from the set of all results for a given system size, these subsets are created from the set of unique evolutions for a given system size. Similarly, the set of the evolutions of the most robust and unique configurations is created to establish correlations between

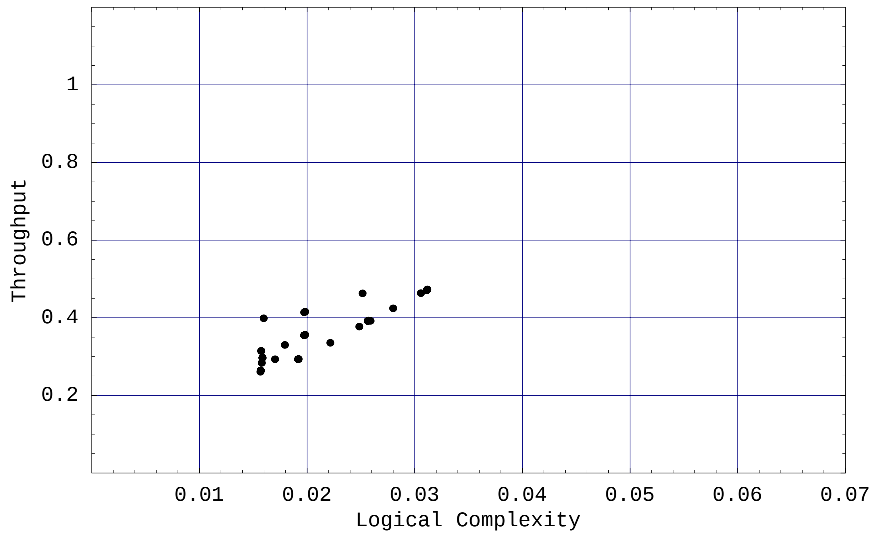

Figure 7.24: The distribution of evolutions unique to 2-2-2 size systems appears similar to the distribution of all evolutions because 82 percent of evolutions are unique.

complexity and throughput and to identify trends applicable in an optimal search.

The two-dimensional distribution of non-halting evolutions unique to 2-2-2 size systems is presented in Figure 7.24. This distribution does not look significantly different than the distribution of all non-halting evolutions shown in Figure 7.21(a). Much of the similarity is a result of the limited number of mimics with respect to carriages - there are just 12 mimics out of a total of 66 evolutions. However, the frequency landscapes are different, with many of the low complexity/low throughput evolutions present in Figure 7.22(a) for the set of all non-halting evolutions absent in the set of unique non-halting evolutions, which is presented in Figure 7.25.

The absolute values of the differences are evident in a comparison of the histograms with respect to logical complexity and throughput, shown in Figure 7.26. There are 54 unique non-halting 2-2-2 evolutions, meaning 12 of the 66 2-2-2 evolutions are equivalent to smaller systems. All 12 evolutions are low logical complexity and low throughput, but represent a mixture of complete and incomplete evolutions.

Perhaps the most direct means of visualizing the differences between the distributions of all non-halting evolutions and the subset of unique evolutions is with histograms of the set of the 12 mimics, shown in Figure 7.27.

In this figure, we see without question that the low throughput/low logical complexity evolutions are mimics. This result should be no great surprise - the only smaller system size that can be emulated by a 2-2-2 system (with respect to carriages) is a 2-1-2 system, which has an equivalent logical complexity and throughput to a 1-1-1 system and only exists at the lowest logical complexity and throughput combination. Mimicry at the minimal logical complexity

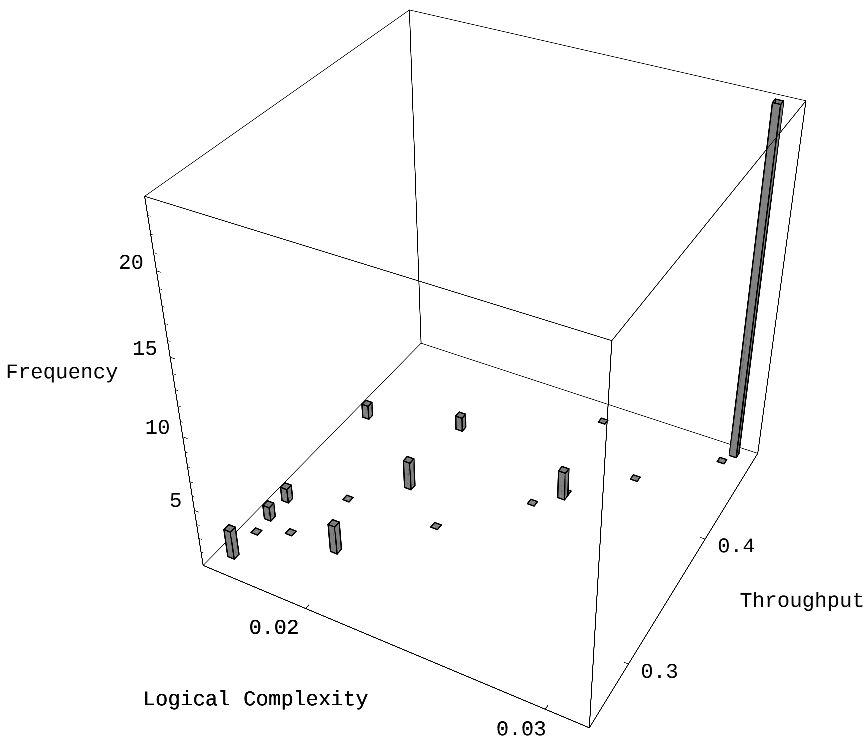

Figure 7.25: The frequency landscape of unique 2-2-2 evolutions reveals that most of the mimics have low complexity and low throughput.

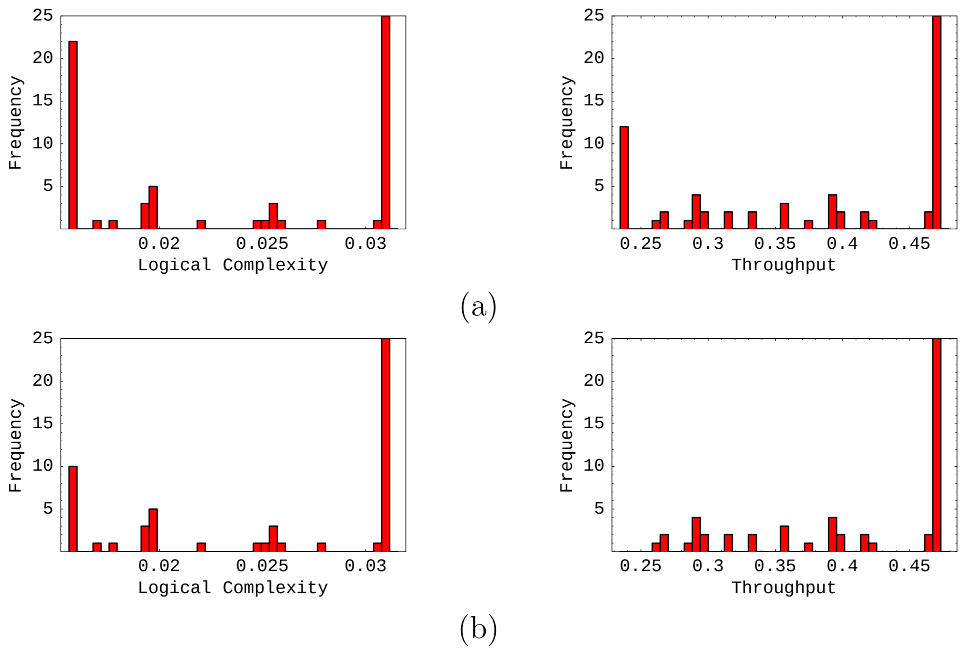

Figure 7.26: The histograms of cross sections of logical complexity and throughput for (a) all non-halting evolutions with mimics included and (b) all non-halting, unique evolutions. The number of high logical complexity/high throughput evolutions remains constant, but the minimal logical complexity/throughput evolutions are not present in the set of unique evolutions.

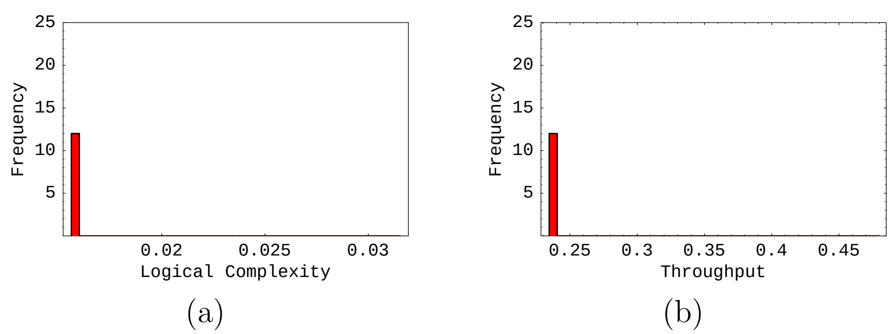

Figure 7.27: The histograms for the evolutions that mimic smaller systems with respect to (a) logical complexity and (b) throughput. Mimics tend to have lower logical complexity and throughput, just like smaller systems.

and throughput is not always the case however. For larger size systems with greater than two carriages, mimicry of more complex dynamics is possible although the simplest evolutions will always be in the set of mimics (there always exists at least one physical configuration in which a single shaft is connected to the magazine corresponding to the single available item type).

Because the logical complexity and throughput are minimal for all mimics, the mean logical complexity and throughput for the set of unique, non-halting evolutions are both greater than for the set of all non-halting evolutions. The mean logical complexity for the unique evolutions is 0.0252 and 0.0235 for all non-halting evolutions while the mean throughput is 0.406 for the unique evolutions and 0.375 for all non-halting evolutions. A higher mean logical complexity and throughput for unique evolutions suggests that increases in the number of carriages result in the ability to support additional complexity in behavior. Additionally, since mean throughput also increases, we can again suggest a correlation between logical complexity and throughput - greater complexity supports greater performance. Once again, these conclusions are tentative, since they are based on small systems with few evolutions. However, they identify trends to examine with the results of larger systems.

The set of unique non-halting evolutions contains incomplete evolutions that may be considered as misrepresenting the true nature of the relationship between complexity and performance. Removal of these evolutions affects the distribution of evolutions and therefore the mean logical complexity and throughput. The distribution of the 34 unique, complete 2-2-2 evolutions is presented in Figure 7.28. The distribution appears identical to the distribution of complete evolutions shown in Figure 7.21(b). In this case, the similarity results from a difference of only six evolutions. As a result, the histograms shown in Figure 7.29 are nearly identical to those in Figure 7.23(b). However, since the six incomplete evolutions have the lowest logical complexity and throughput from the set of unique, non-halting evolutions, the mean logical complexity and throughput for the set of unique, complete evolutions are slightly higher (0.0283 vs. 0.0264 for logical complexity and 0.431 vs. 0.402 for throughput). This difference verifies what we observe in a comparison of histograms - that unique, incomplete evolutions have a lower aggregate logical complexity and throughput compared to the entire set of unique, non-halting evolutions.

The set of unique, robust evolutions is a subset of the unique, complete evolution set and contains only the evolutions of configurations that completely deliver all queue distributions. The two-dimensional distribution of the evolutions of robust configurations is shown in Figure 7.30. The 18 evolutions from the three configurations that populate this set share the same logical complexity of 0.0311 and throughput of 0.471, which correspond to the maximum values for the set of complete evolutions. Since all logical complexities and throughputs from this set are identical and maximal, the mean logical complexity and throughput are greater than the means for both unique, non-halting evolutions and unique, complete evolutions. As with the set of non-unique, robust configurations, it appears as if a correlation between logical complexity and throughput is possible, based on the distinction of robustness. If we associate adaptability with robustness, then the most adaptable configurations

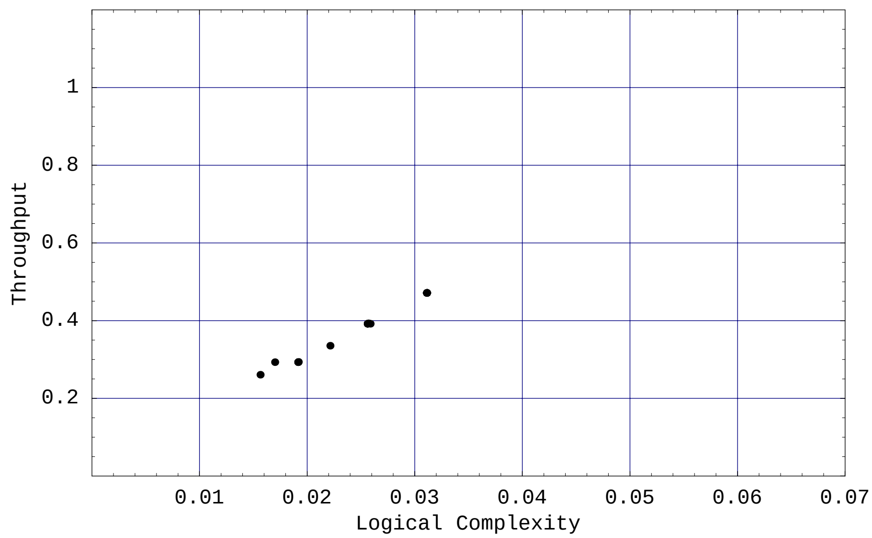

Figure 7.28: The two-dimensional distribution of unique, complete 2-2-2 evolutions is nearly identical to the distribution of all complete evolutions primarily because the sets differ by only 6 evolutions. The 6 evolutions have low logical complexity and throughput, resulting in a greater mean logical complexity and throughput for the unique evolutions.

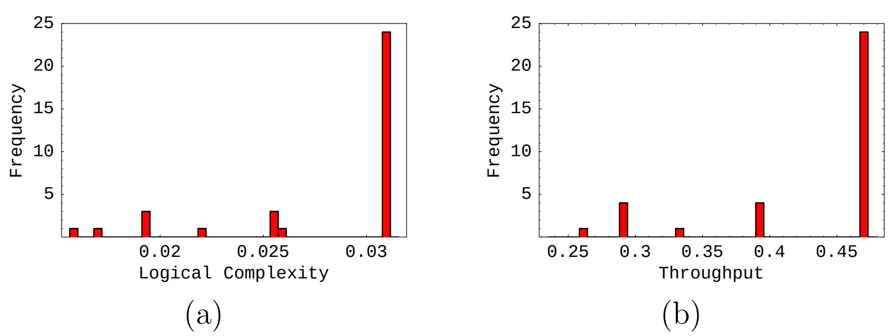

Figure 7.29: The histograms for unique, complete 2-2-2 evolutions reveals the logical complexity and throughput of the 6 missing evolutions included in the histograms for all complete evolutions

are also the most logically complex and exhibit the greatest throughput. The distinction of uniqueness shows that increases in system size are directly related to adaptability.

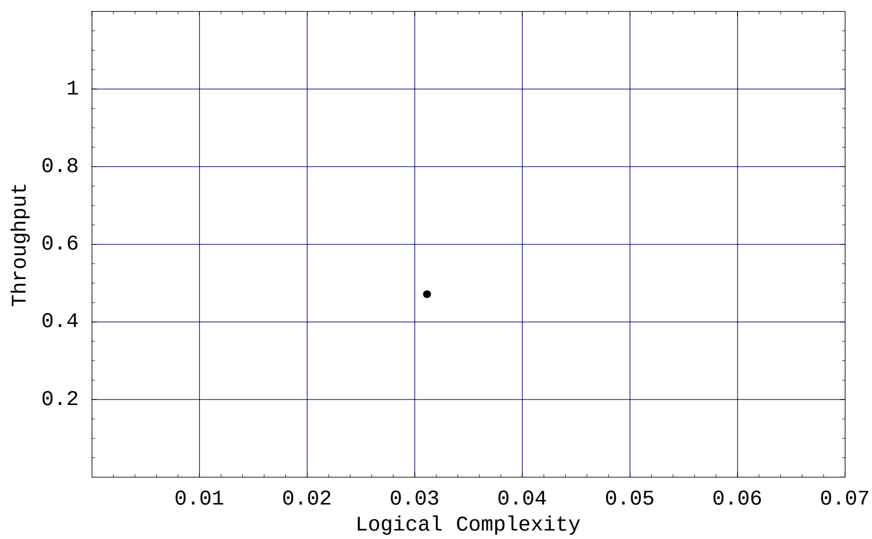

Figure 7.30: The two-dimensional distribution of the evolutions of the unique and robust 2-2-2 configurations. The members of this set are unique and complete for all queue distributions considered.

The three configurations of this set are 111, 127, and 255. Example evolutions of 255 have already been presented in Figure 7.15 for (0-100) and (60-40) queue distributions. The remaining evolutions in this set share the same characteristics as these examples: regardless of the queue distribution, the carriages become phase-locked with a phase-lag that is sufficient to result in effectively independent carriages. The configurations also share another trait. In order to accomplish complete delivery for all queue distributions and to be unique to a given system size, the shafts must be connected to all magazines. The incidence matrices of these configurations, shown in Figure 7.31, reveals that 111, 127, and 255 are the only configurations (without trivial repetitions) that have this characteristic. The connectivity to the queues may affect the order of the evolutions, but the net effect is the same: all items are delivered for all input streams and a phase lag between carriages always occurs when the distribution of item types is unequal, resulting in resource sharing.

The complete connectivity of shafts and magazines in unique, robust configurations results in all non-zero entries in the QM matrix, meaning there is always at least one path from each queue to each magazine. Completely non-zero QM matrices are a characteristic of configurations that always have complete evolutions but do not necessarily correspond to uniqueness for all queue distributions.

The incidence matrices for configurations 246 and 247 are presented in Figure 7.32. The non-zero QM matrices for these configurations are a result of the complete connectivity of shafts and queues and, while they guarantee complete delivery, the incomplete shaft-

$$\text{(a)}\quad (SQ) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$$

$$\text{(b)}\quad (SQ) = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 1 & 1 \\ 2 & 2 \end{pmatrix}$$

$$\text{(c)}\quad (SQ) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 2 & 2 \\ 2 & 2 \end{pmatrix}$$

Figure 7.31: The incidence matrices for the unique, robust configurations (a) 111, (b) 127, and (c) 255. All have complete shaft-magazine connectivity and the different configurations reflect the possible variations of shaft-queue connectivity.

magazine connectivity assures that the configurations will mimic a smaller system for at least one queue distribution. Configuration 246 mimics a 2-1-1 (or 2-1-2) system when only a single item type exists and configuration 247 mimics a 2-1-1 system when all items are bound for the first magazine.

$$\text{(a)}\quad (SQ) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$$

$$\text{(b)}\quad (SQ) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 1 & 2 \\ 1 & 2 \end{pmatrix}$$

Figure 7.32: The incidence matrices for configurations (a) 246 and (b) 247 show complete connectivity between shafts and queues and non-zero QM matrices. These configurations always result in complete evolutions, but are not unique for some queue distributions.

The incidence matrices for configuration 119 are presented in Figure 7.33 and show a non-zero QM matrix. However, neither the shaft-queue matrix nor the shaft magazine matrix is completely connected. The first shaft is always used regardless of the queue distribution because it is connected to both queues and magazines, but the second shaft is only used when items bound to the second magazine (in the second queue) are present. Therefore, when all items are bound for the first magazine, the first shaft is not utilized and the system mimics a 2-1-1 system.

Uniqueness of a configuration (on a per evolution basis) is a function of the queue distribution. If the shafts and magazines of a configuration are connected sufficiently with respect to the queue distribution (i.e. each shaft must access at least one of the magazines corresponding to the item types in the queues, i.e the sub-matrix created from the SM matrix consisting of the columns corresponding to magazines with present item types has no all-zero rows), then all shafts are utilized and the configuration is unique *with respect to the queue*

$$(SQ) = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$$

Figure 7.33: The incidence matrices for configuration 119 shows a non-zero QM matrix but incomplete connectivity between shafts and magazines and shafts and queues. 119 will always be complete, but not always unique.

distribution.

It is possible for an incomplete evolution to result from a configuration with a sufficiently connected SM matrix depending on the connectivity of shafts and queues. Therefore, a configuration with a complete evolution for a given queue distribution has non-zero rows in the SM sub-matrix of all shafts and those magazines corresponding to present item types in the queue distribution and an SQ matrix such that the QM sub-matrix for all queues and the relevant magazines is non-zero.

Robust configurations (always complete) will therefore have non-zero QM matrices, but can have incomplete connectivity between shafts and magazines. For uniqueness for a given queue distribution, a configuration must have complete connectivity in the SM matrix for the magazines involved in the queue distribution. So for robust uniqueness (always unique), it follows that the SM matrix must be completely connected for all magazines.

7.4 Larger Systems

Subtle differences occur to the correlation between the logical complexity and throughput with increases in system size that are evident in both the sets of all evolutions and unique evolutions. To exhaustively analyze all 59 conventional sets in terms of the subsets described for 2-2-2 size systems is of little practical value. Instead, we will investigate particular systems that reveal relationships between complexity and performance unseen in 2-2-2 evolutions and will identify the dependence of these relationships to the relative ratios of queues, shafts, and magazines. We also explore the larger systems for evidence that supports the tentative conclusions regarding logical complexity and throughput in 2-2-2 systems.

7.4.1 2-4-2 Systems

To obtain an idea of the effect of raising the number of shafts in relation to the number of queues and magazines, we consider the set of 2-4-2 evolutions. Figure 7.34 presents the two-dimensional distribution of the entire set of 1217 2-4-2 evolutions, less the 31 terminated evolutions. With a considerably larger number of evolutions, the distribution and clustering of evolutions is more meaningful. The clustering of evolutions for this set is apparent in the frequency landscape in Figure 7.35. The distributions reveal that the near linear right bound that exists for the set of 2-2-2 evolutions is also present for this set. The left and right boundaries also diverge from minimal values of logical complexity and throughput, although the slope of the left bound is not as steep as that for 2-2-2 size systems. The upper bound in Figure 7.34 is significantly more defined than the upper bound for 2-2-2 systems and is approximately a horizontal line, indicating a maximum throughput over a wide range of logical complexities. The absolute value of the maximum throughput is larger than that for the 2-2-2 systems, as is the maximum logical complexity. However, the minimal logical complexity/minimal throughput value is identical to the minima for the 2-2-2 systems. The frequency landscape for the 2-4-2 evolutions reveals peaks at the same relative low logical complexity/low throughput and high logical complexity/high throughput values as for 2-2-2 evolutions, although the most significant cluster of evolutions in the 2-4-2 evolution set occurs at moderate to low values of logical complexity and throughput, a relative space that is unoccupied in by 2-2-2 evolutions.

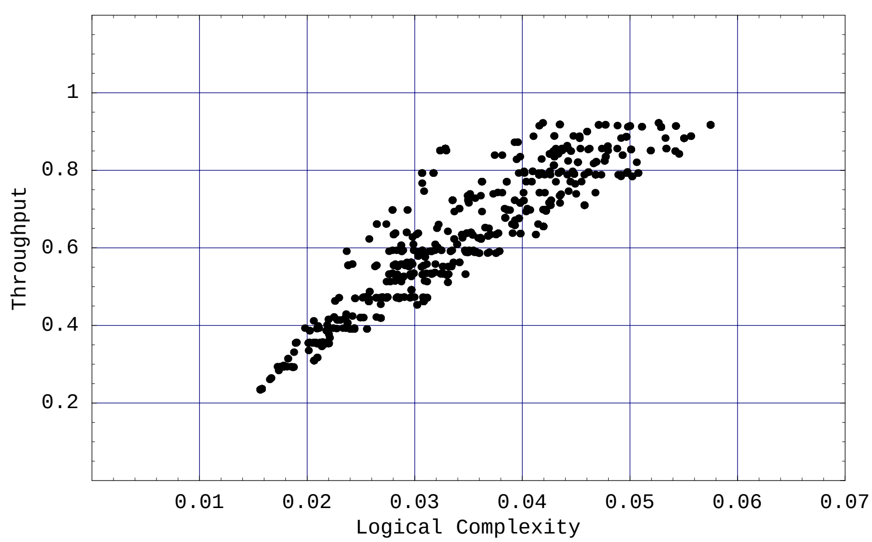

Figure 7.34: The two-dimensional distribution of 2-4-2 evolutions represents a much larger set than for 2-2-2 evolutions. The near linear right boundary is still apparent and an upper boundary is more defined. The slope of the left boundary is not vertical.

The presence of significantly more evolutions enables a more detailed analysis of the characteristics of the right boundary. In 2-2-2 systems, we saw that along the right boundary, all evolutions are mixtures of the maximum and minimum logical complexity/throughput combinations. Position along the boundary is defined by the relative proportion of each evolution type. Evolutions along the boundary experience the halting of one or more carriages in the course of the evolution and maximum state transitions when carriages are in operation. When a carriage halts, the diversity of states decreases and there is the potential for fewer state transitions per system cycle, which tends to decrease logical complexity. Halting carriages also have the effect of increasing the number of trips required by the remaining carriages to deliver all items, increasing the total evolution time and decreasing the

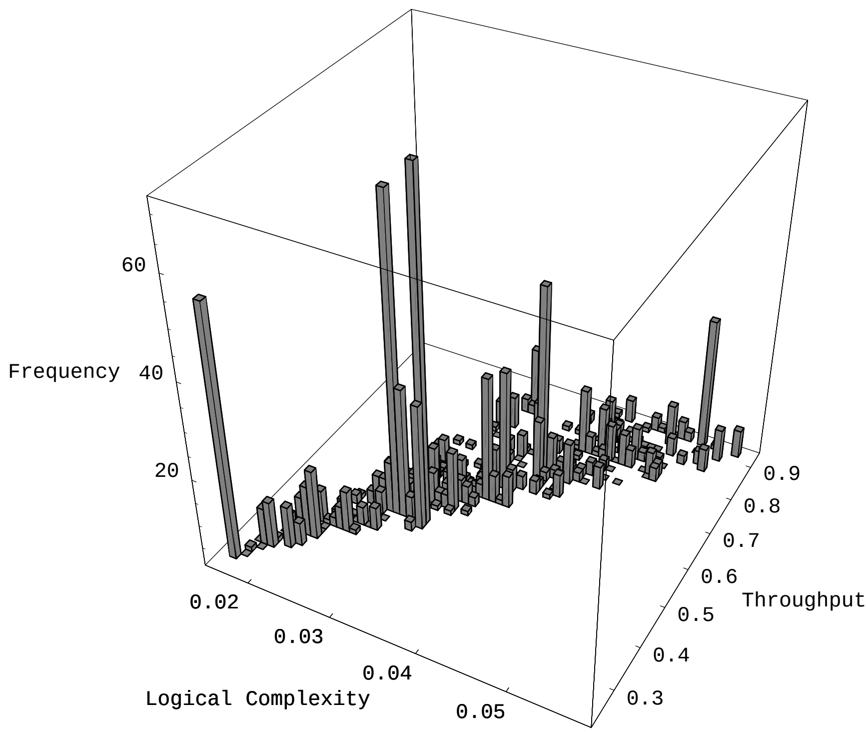

Figure 7.35: The three-dimensional frequency landscape has the same peaks at the minimum and maximum logical complexity/throughput combinations as the 2-2-2 evolution set, but is dominated by mid valued logical complexity and throughput evolutions. Note other peaks, like the maximum throughput/mid logical complexity peak.

throughput for the same number of items transported. The interdependence of these effects result in a proportionate change in logical complexity and throughput and the near linear relationship observed. All carriages in these evolutions are interdependent, either directly or indirectly, resulting in phase lags between all carriages. A maximum number of phase lags equal to the number of carriages results in the maximum possible logical complexity (assuming synchronized carriage evolutions). The maximum number of phase lags therefore defines the maximum logical complexity while the points in an evolution at which halting of carriages occurs defines the slope of the boundary. There is no reason the logical connectivity is limited to this boundary line for systems with non-deterministic operation times, magazines distributed at different levels, or for different relative ratios of deterministic operation times. The number of states visited per evolution, or at least for one system cycle, for one or any number of these conditions can therefore change the maximum logical complexity. However, more states per cycle implies the possibility of more carriage interactions, which tends to decrease throughput.

The effects of carriage halting with maximum possible logical complexity is best illustrated through examining evolutions at different points on the boundary. Figure 7.36 shows the position in the two dimensional evolution distribution of configuration 32347 with a (80-20) queue distribution and configuration 62831 with a (60-40) queue distribution.

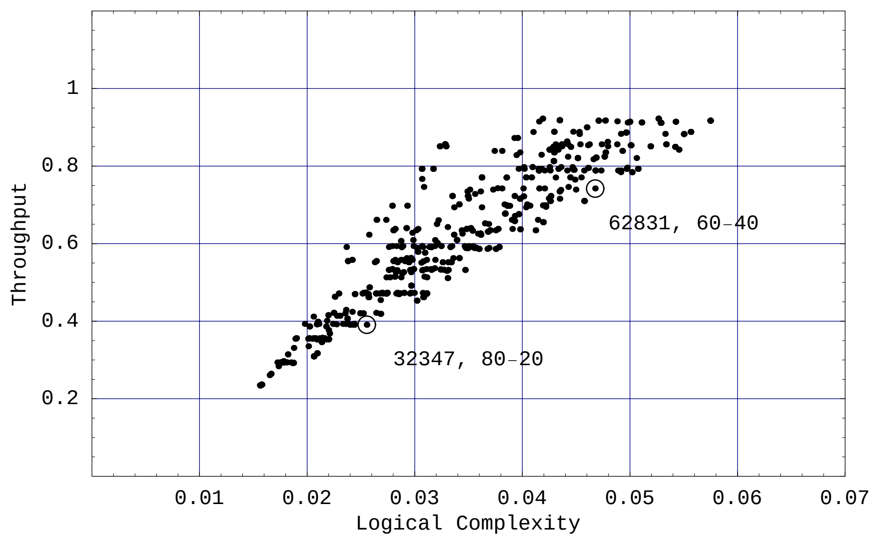

Figure 7.36: The positions of two example 2-4-2 evolutions, 32347 (80-20) and 62831 (60-40), on the right boundary line.

The compressed carriage evolution history for configuration 32347 with a (80-20) queue distribution is presented in Figure 7.37. Additional detail is provided in Figure 7.38, showing the temporal carriage histories in the beginning of the evolution. The evolution begins with the introduction of phase lags between all carriage cycles, a characteristic shared by the evolution with maximum logical complexity and evident in the early part of Figure 7.38 approximately up to step 1100. The logical complexity is therefore maximal for the number of carriages operating in the deterministic environment. Early in evolution however, the first and second carriages halt, and the fourth carriage soon follows. The operational periods of the carriages is readily visible in Figure 7.39, where a flat line indicates a halted carriage.

In the period when the third and fourth carriages are the remaining carriages in operation, the phase lag between them is still present, evident in the latter portion of Figure 7.38. The presence of the phase lag maintains the maximum possible logical complexity for the number of carriages in operation. The *absolute* value of the logical complexity is lower with only two carriages however, and the average logical complexity of the entire evolution decreases. The logical complexity eventually goes to a minimum value corresponding to the lone operational carriage.

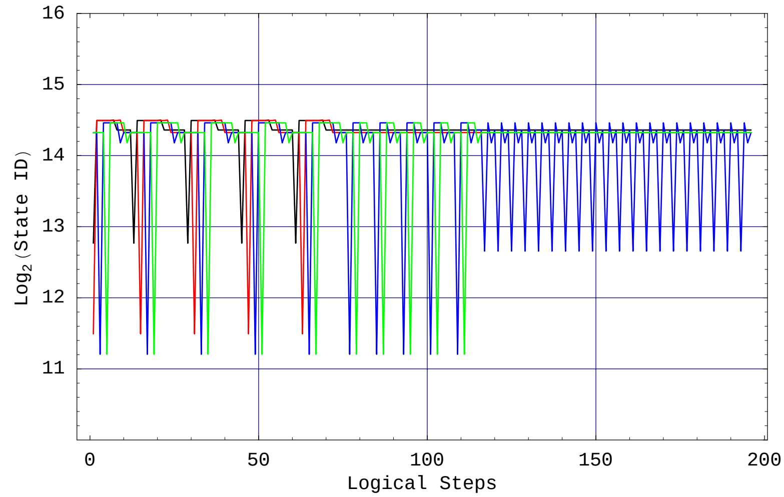

Figure 7.37: The compressed carriage histories of evolution 32347 (80-20). The evolution undergoes several stages, corresponding to the halting of carriages. Note that throughout the evolution, carriages are always phase lagged with each other, resulting in maximum logical complexity for the number of carriages in operation - a requirement for existing on the right boundary.

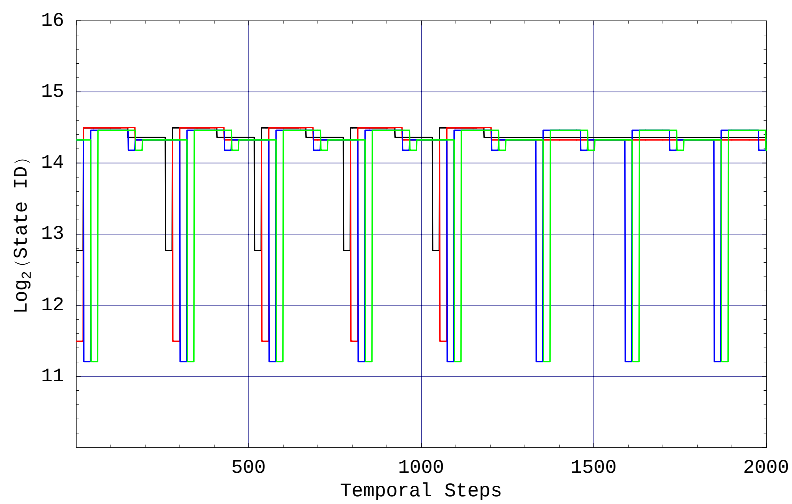

Figure 7.38: Detail of the temporal carriage histories of evolution 32347 (80-20) from 0 to 2000 time steps. The phase lag between carriages is evident as is the difference between absolute logical connectivities for four carriages in operation (up to step 1100) and two carriages in operation.

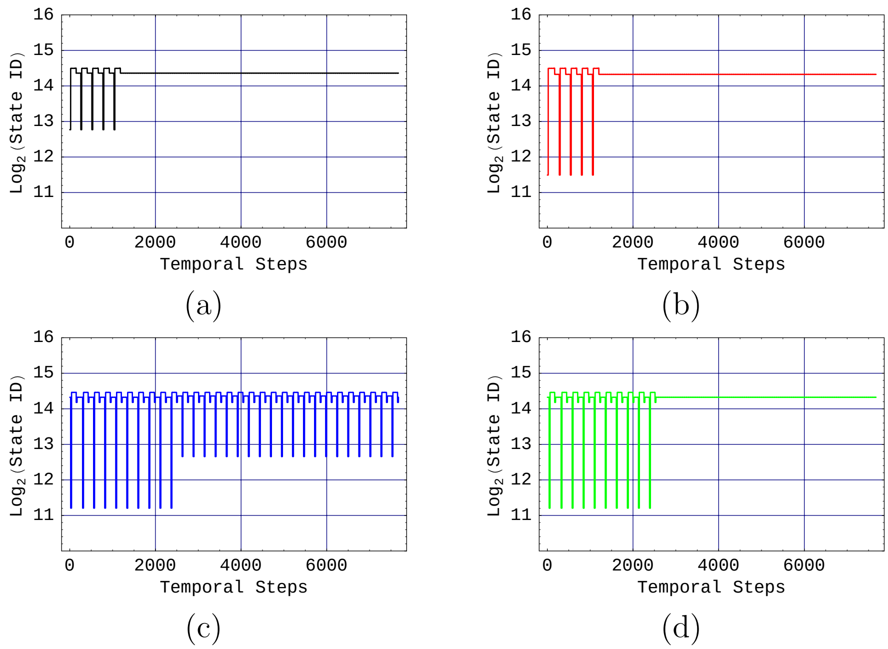

Figure 7.39: The individual temporal evolution histories for (a) carriage 1, (b) carriage 2, (c) carriage 3, and (d) carriage 4 for configuration 32347 with an (80-20) queue distribution. Carriages 1 and 2 halt early in the evolution, followed soon by carriage 4. Carriage 3 completes the delivery of all items.

The compressed carriage state histories for evolution 62831 (60-40) is presented in Figure 7.40 and shows a similar evolution as 32347 (80-20). The phase lags between carriages are still present, evident in the detailed temporal carriage state histories for the first 2000 time steps, shown in Figure 7.41, although the actual trajectories and order of carriages is different. The phase lags again result in the greatest possible logical complexity for the number of carriages in operation and raise the average evolution logical complexity. Figure 7.42, which shows the individual temporal carriage state histories, indicates that the carriages halt much later in this evolution than for 32347 (80-20). Since more of the evolution has greater local logical complexity because more carriages remain in operation longer, the average logical complexity is greater. Additionally, more carriages in operation implies shorter evolution times and therefore greater throughput.

The logical complexity of an evolution along the right boundary is therefore simply a weighted average of the "local" logical complexities of periods corresponding to the number of operational carriages. For evolutions with carriages that halt early, lower logical complexities dominate the average. Throughput is lower in these evolutions because fewer carriages require more time to deliver a given set of items. The opposite is true for evolutions in which carriages halt later.

The linear boundary defining maximum logical complexity stems from this characteristic.

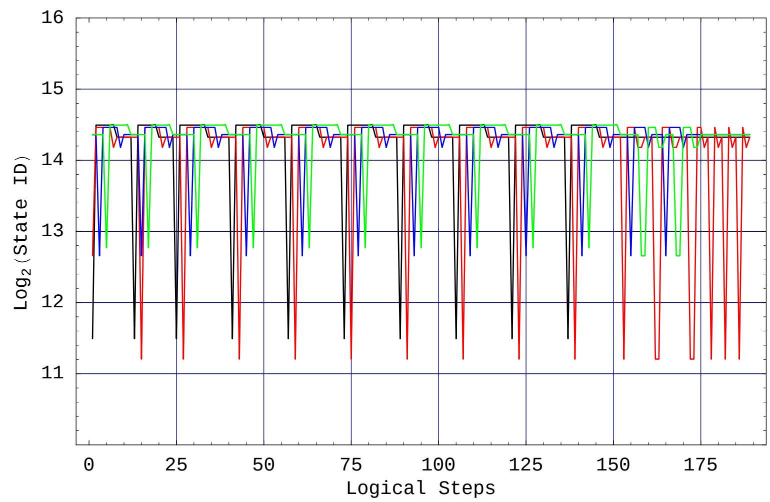

Figure 7.40: The compressed carriage histories of evolution 62831 2-4-2 (60-40) are similar to those of evolution 32347 2-4-2 (80-20). The phase lags between every carriage result in maximum logical complexity for the number of carriages in operation. However, the carriages remain in operation longer for 62831 (60-40), resulting in a higher average logical complexity.

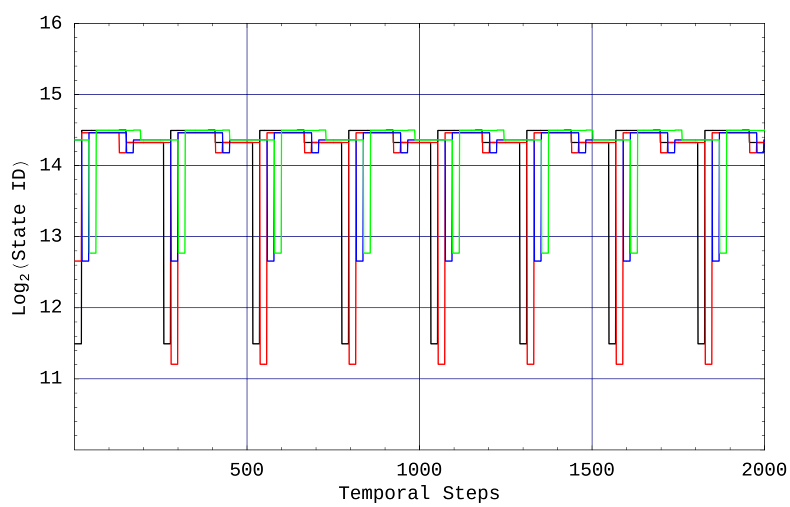

Figure 7.41: The detail of the temporal carriage histories of evolution 62831 (60-40) from 0 to 2000 time steps illustrates the presence of the phase lags between all carriages.

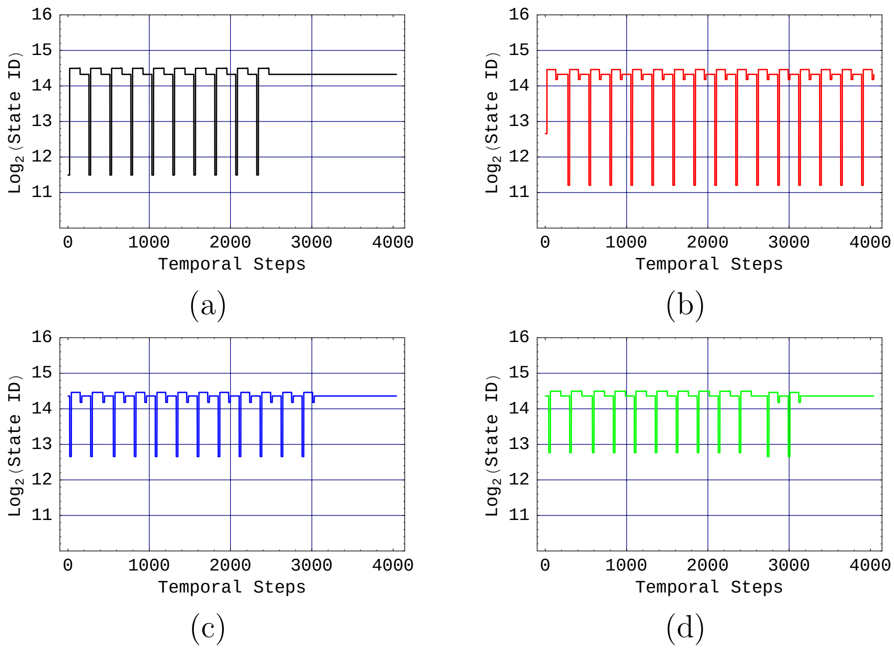

Figure 7.42: The individual temporal evolution histories for (a) carriage 1, (b) carriage 2, (c) carriage 3, and (d) carriage 4 for configuration 62831 with a (60-40) queue distribution. Carriages halt later in this evolution than for 32347 (80-20), raising the average logical complexity.

Throughput is directly related to the number of carriages in operation. This relationship is observed in a comparison of the minimum and maximum throughputs for any evolution set. The fractions of the evolution corresponding to possible numbers of carriages in simultaneous operation therefore affects the logical complexity and the throughput in direct proportion, as long as the increase in the number of logical steps per additional operational carriage is constant.

The effects of the averaging of local logical complexities and the direct relationship of throughput and the number of operational carriages that defines the right boundary when phase lags exist between all carriages are not restricted to the maximum logical complexity boundary. These effects are extendable throughout the logical complexity/throughput field and can be used as a tool for making comparisons between evolutions and defining equivalent configurations. To illustrate the extension of these effects, we look at the left boundary line - the minimum logical complexity boundary.

The discrepancy between the slope of the left boundary of the distribution of evolutions for 2-2-2 and 2-4-2 size systems is attributable to the relative ratio of location types. In 2-2-2 size systems, the independent, phase locked circuits found at the top of the left boundary in Figure 7.21(a), represented by configurations with minimal connectivity such as 102, are possible. However, in 2-4-2 size systems, no truly independent circuit is possible because of the ratio of shafts to queues and magazines. At least two shafts are always sharing at least one queue or magazine in these configurations by definition of a valid configuration. With no independent circuits possible, a minimum of one phase lag exists that results in additional logical complexity. The minimum of a single phase lag, whether embodied as two sets of synchronous carriages with a phase lag between sets or a set of three synchronous carriages phase lagged with a lone carriage, defines the left boundary along with the proportional relationship between throughput and the fraction of carriages operating simultaneously. Since logical complexity will always be greater in an evolution with a single phase lag than for a simple evolution of a lone carriage, the slope of this left boundary is therefore positive.

An example evolution that exists on the left boundary is configuration 57215 with a (60-40) queue distribution. The compressed carriage state histories for this evolution are presented in Figure 7.43. This evolution is dominated by a repetitive pattern that persists for approximately 80 logical steps. A detailed view of the pattern is visible in Figure 7.44, which shows only a single phase lag. The first and second carriages change states simultaneously, as do the third and fourth carriages. However, the unavoidable sharing of queues initiates a phase lag between the sets of carriages.

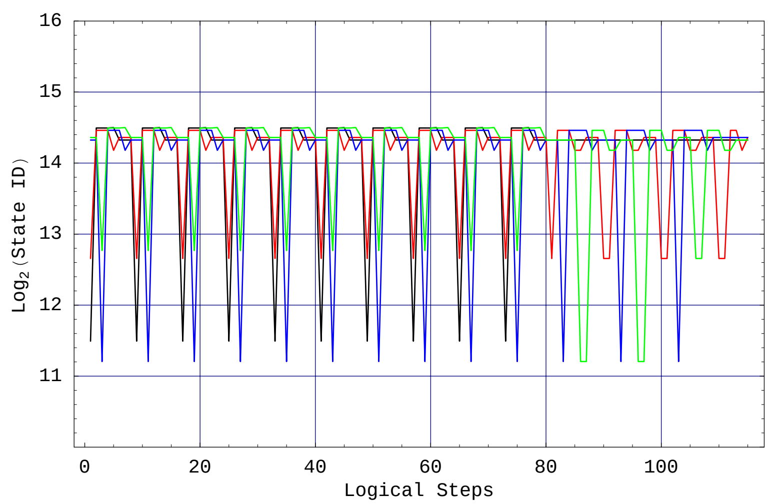

Figure 7.43: The compressed carriage histories of evolution 57215 2-4-2 (60-40). Carriages 1 and 2 are phase lagged with carriages 3 and 4, resulting in a single system phase lag.

Evolution 57215 (60-40) is the evolution from the set of 2-4-2 systems with the greatest possible throughput with a single phase lag. However, this evolution has a lower throughput than the maximum possible, or achieved, in a 2-4-2 system. The less-than-maximal throughput results not from the effects of changes in the number of states per system cycle or the breakdown of the dominant pattern at the end of Figure 7.43, but because the first carriage halts approximately $\frac{3}{4}$ of the way through the evolution. While a maximum throughput is theoretically possible, the position of evolution 57215 (60-40) illustrates that the effects of halting carriages evident on the right, maximum logical complexity boundary are also ap-

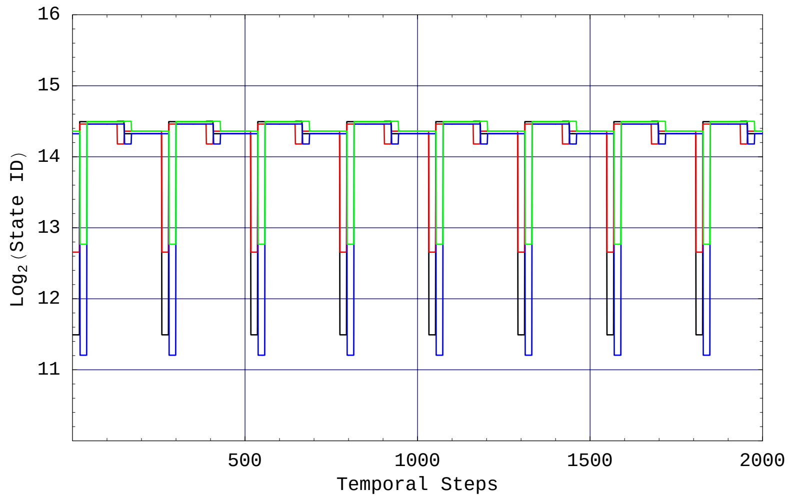

Figure 7.44: The detail of the temporal carriage histories of evolution 57215 2-4-2 (60-40) for the first 200 time steps shows the single phase lag that exists between the two carriage sets.

plicable along the left, minimum logical complexity boundary. As carriages halt earlier, the logical complexity drops in proportion with the throughput, creating a linear relationship corresponding to a specific number of phase lags, assuming the evolution always maintains a constant number of effective phase lags between carriage sets.

In evolution 57215 (60-40) and similar evolutions, where a single phase lag is present between groups of synchronous carriages, the collective system state comprised of the individual carriage states in the group is equivalent to a single carriage with respect to logical complexity. Logically, the evolution can therefore be thought of as emulating a simpler system with only two carriages[4]. Since the number of phase lags is related to the logical complexity and emulated system size, the left and right boundaries defined by the minimum and maximum phase lags also indicate the emulated system size. The right boundary corresponds to self-emulation, with all possible carriages present because all carriages are required to obtain the maximum number of phase lags associated with this boundary. The left boundary is always the set of evolutions that emulate evolutions with the smallest number of carriage groups possible for that system size, depending on the number of magazines and queues.

Emulation is further illustrated with a comparison of the logical complexity of unique 2-2-2 evolutions and 2-4-2 evolutions that emulate the logical complexity of two carriage configurations. The maximum logical complexity for a 2-2-2 system, corresponding to maximum throughput, is 0.0311, while the logical complexity for evolution 57215 (60-40), the 2-4-2

[4]By definition, this equivalence is not the same as mimicry. In mimicry, a carriage is completely absent while an emulation of a smaller system requires the participation of more carriages than present in the mimicked system. A system that mimics a system with a carriages can therefore also emulate a system with b carriages as long as $a > b$.

system with the greatest throughput that emulates a 2-2-2 system is 0.0328. The difference between these values is approximately 5%, which can be attributed to the effect of the halting carriage in 57215. A comparison of the throughputs for these evolutions reveals that emulation, not mimicry, is at work. The throughput of evolution 57215 (60-40) is 0.856 while the throughput of the 2-2-2 evolution is 0.473. The throughput of the 2-4-2 system is 1.81 times the emulated 2-2-2 system, close to the doubling expected from twice the number of *operational* carriages. The difference is again attributed to the halting of the first carriage in 57215.

7.4.2 Qualitative Characterizations

Based on the relationship between phase lags, the occurrence of halting, and emulations, we can draw theoretical minimum and maximum logical complexity boundaries drawn from characteristics of evolutions near these lines. The temporal number of steps required for one carriage to complete one carriage cycle (one round trip from a queue to a magazine) is 258. The system cycle time in a deterministic system is therefore also 258 time steps regardless of the number of carriages, as long as the product of the number of carriages and the time delay between carriages is less than the carriage cycle time. However, the number of logical steps increases linearly with the number of phase lags. For the minimum of a single phase lag for 2-4-2 systems, the number of logical steps is 8, illustrated in a section of the compressed carriage history for evolution 57125 (60-40) in Figure 7.45(a), while 16 logical steps occur per system cycle in an evolution with the maximum of three phase lags as in Figure 7.45(b).

If the pattern shown in Figure 7.45(a) continues indefinitely, and we omit transients, the theoretical logical complexity is 0.0310, very close to the actual values found for 255 2-2-2 (any queue distribution) and 57215 2-4-2 (60-40). The theoretical logical complexity for an evolution with three phase lags and 16 logical steps per system cycle is 0.0620, which is simply double the theoretical logical complexity for a single phase lag because there are twice the number of logical steps for the same number of temporal steps. The theoretical throughput when each pattern is continued indefinitely is 0.930 regardless of the number of phase lags because four carriages transport four items in the same time in both evolutions $(\frac{4items}{258seconds} \cdot \frac{60s}{min})$. Once again, this is only slightly greater than the maximum throughput of 0.918 from the set of complete evolutions.

Knowing that the halting of carriages at different times throughout an evolution affects the logical complexity and throughput linearly, we can construct lines between the maximum throughput and logical complexity values corresponding to the number of phase lags and the minimum logical complexity and throughput values corresponding to an evolution with a single operational carriage transporting all items. The minimum theoretical values for logical complexity and throughput, 0.0155 and 0.232, respectively, are very near the actual minimum logical complexity and throughput of 0.0156 and 0.235. We can also construct a maximum throughput boundary connecting the maximum theoretical throughputs for

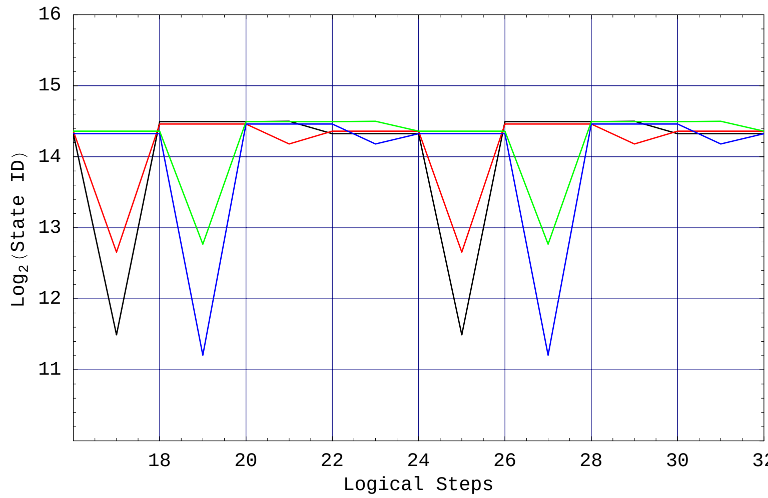

(a)

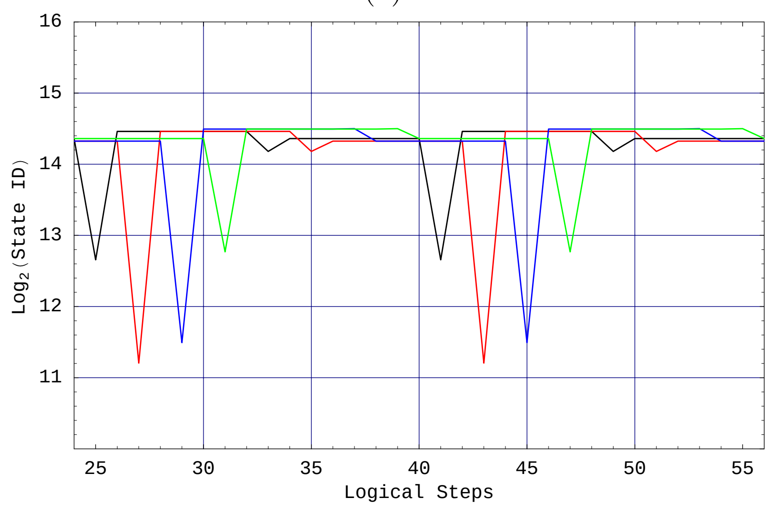

(b)

Figure 7.45: A minimum of a single phase lag in a 2-4-2 evolution results in 8 logical steps per system cycle, shown in a detail of the compressed evolution of 57215 (60-40) in (a). 16 logical steps per system cycle occur when the number of phase lags is at a maximum of three, as shown in the compressed carriage history in (b), using configuration 32767 (60-40) as an example.

evolutions corresponding to the minimum and maximum number of phase lags knowing that more than four items can not be carried in the smallest system cycle length with only four carriages in the system. The resulting boundaries for the 2-4-2 system are shown in Figure 7.46 for the set of complete evolutions, which reveals all evolutions are contained within the theoretical bounds.

The relationship between the number of phase lags and the linear distribution of evolutions

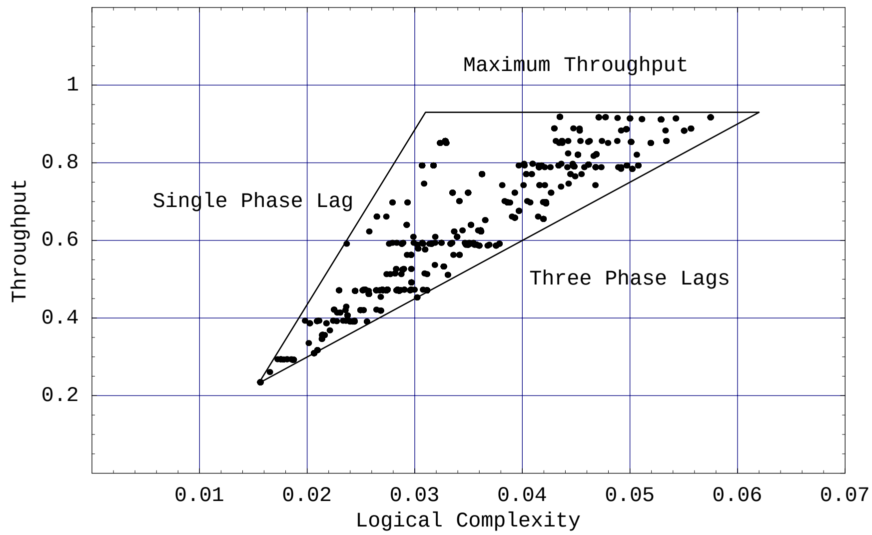

Figure 7.46: All complete evolutions (and non-halting evolutions) are contained within the bounds of the theoretical limits on logical complexity and throughput, which are determined based on the extreme number of possible phase lags and the relationship between the fraction of halting carriages and throughput and logical complexity.

is extendable to all possible levels of emulation in addition to those corresponding to the minimum and maximum number of carriages possible. In addition to emulation of 2-2-2 systems, emulation of 2-3-2 systems is also possible by 2-4-2 systems and corresponds to an evolution with two phase lags. An example of an evolution with two phase lags is shown in Figure 7.47, which shows the detailed compressed carriage state histories for 24 logical steps for evolution 24575 (40-60).

If the two phase lag pattern is repeated indefinitely, there are 12 logical steps per 258 temporal steps, resulting in a maximum theoretical logical complexity of 0.0465, directly between the maximum theoretical logical complexities of 2-4-2 and 2-2-2 emulated evolutions. We can therefore include a line representing 2-3-2 emulations in Figure 7.48, position along which is determined based on the average number of operational carriages over the course of an evolution.

Conceptually, the description of the degree of emulation is not restricted to integer values of the number of carriages. In the same sense as the position along any ray from the minimum logical complexity/throughput value corresponds to the average of the number of carriages in simultaneous operation, the horizontal position of an evolution at a given throughput is the average number of phase lags present and therefore the size of the emulated configuration with respect to the number of carriages. We can therefore think of the logical complexity/throughput space being constructed of an infinite number of rays originating from the minimum theoretical logical complexity/throughput value with slopes between the self-emulated and minimally emulated rays that correspond to real valued average numbers

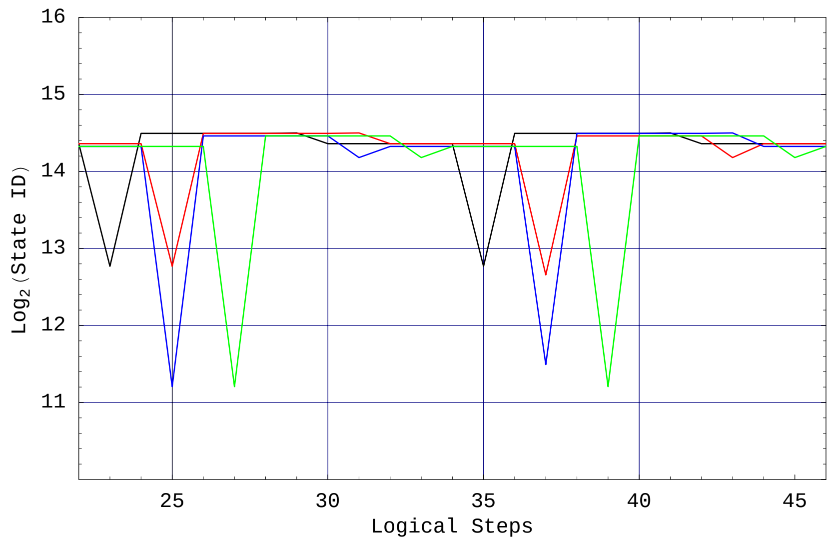

Figure 7.47: Detail of the compressed carriage state histories for evolution 24575 2-4-2 (40-60). Two phase lags result in 12 logical steps per system cycle, which emulates a 2-3-2 configuration.

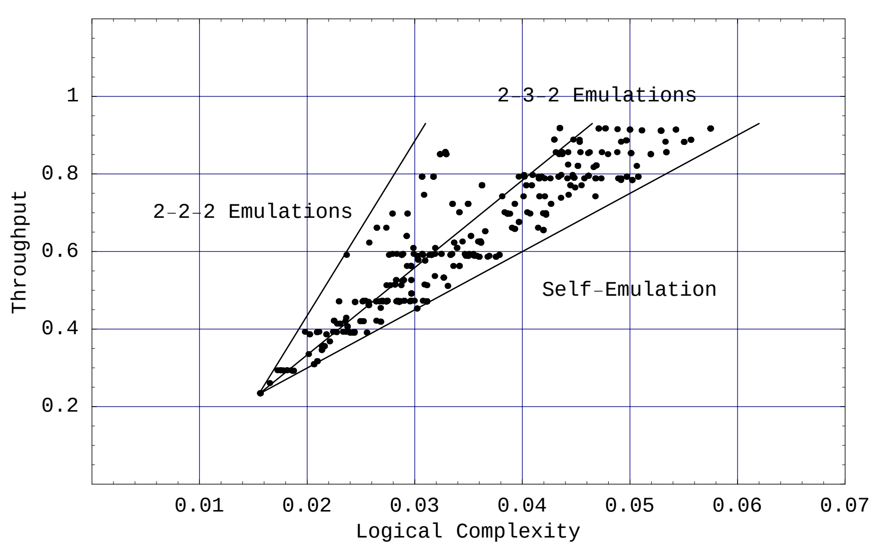

Figure 7.48: Internal boundaries based on the number of emulated carriages are also possible when some carriages are independent and operate without a phase lag, but these carriages share phase lags with other carriages. For the 2-4-2 configurations, 2-3-2 systems are emulated when there two phase lags exist.

of phase lags. An equivalent carriage emulation value is defined as the ratio of the projection of the position of an evolution in the logical complexity/throughput space along a ray originating from the minimum theoretical logical complexity/throughput point onto the maximum theoretical throughput line to the maximum theoretical logical complexity for that system size, described symbolically in Equation 7.1.

$$E_{C_L} = \frac{S\left(R_\mu\left(Cx_{L_i} - Cx_{L_\nu}\right) + Cx_{L_\nu}R_i - Cx_{L_i}R_\nu\right)}{Cx_{L_\mu}\left(R_i - R_\nu\right)} \tag{7.1}$$

where: E_{C_L} is the equivalent carriage emulation
Cx_L is the logical complexity
R is the throughput
S is the number of shafts in the configuration
μ denotes the theoretical maximum value for the entire set
ν denotes the theoretical minimum value for the entire set
i denotes the value of the evolution

By similar triangles, the equivalent carriage emulation value could also be defined as the ratio of an evolution's logical complexity to the maximum logical complexity at the evolution's throughput. As an example, consider evolution 26479 2-4-2 (80-20) with a logical complexity of 0.033 and a throughput of 0.723. The projection of the evolution's value onto the maximum theoretical throughput boundary, illustrated in Figure 7.49, is 0.0411 with a minimum theoretical logical complexity of 0.0155 and a minimum theoretical throughput of 0.232. The ratio of this projection to the maximum theoretical logical complexity for the system of 0.062 results in an equivalent carriage emulation value of 2.65, meaning the 26479 (80-20) evolution effectively has the equivalent logical complexity of an evolution with 1.65 phase lags.

The fact that all lines of constant emulated carriage size originate from the minimum theoretical logical complexity/throughput point raises two apparent contradictions. How can different lines of constant carriage emulation share a common point and how can a line of constant carriage emulation emulate two different numbers of carriages? The answer is related to the definition of emulation and the characterization of evolutions with respect to the average number of operational carriages. In addition to the emulated system size with respect to the number of carriages, the line drawn through an evolution from the minimum theoretical logical complexity/throughput also yields the average number of operational carriages for the evolution. The average number of operational carriages is linearly related to the position of an evolution along a line of constant emulated carriage size with the minimum average number of operational carriages corresponding to the minimum theoretical throughput and the maximum average number of operational carriages corresponding to the maximum theoretical throughput. The average number of operational carriages is expressible as the ratio of the evolution throughput to the maximum theoretical throughput,

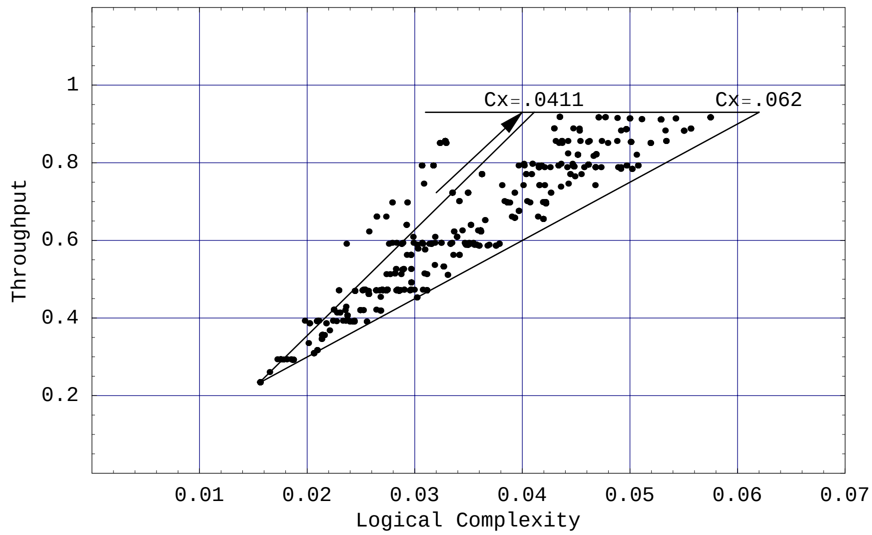

Figure 7.49: Projection of the logical complexity and throughput for evolution 26479 onto the maximum theoretical throughput boundary along the ray originating from the minimum theoretical logical complexity/throughput value provides the equivalent carriage emulation value.

scaled by the number of shafts in the configuration. This result is equivalent to the ratio of the evolution throughput to the minimum theoretical throughput. Symbolically, the relation is presented in Equation 7.2, where E_{O_L} is the logically equivalent average number of operational carriages.

It must be recalled that the definition of emulation is based on logical equivalence and does not describe the actual number of operational carriages in operation. The intersection of lines of constant carriage emulation at the minimum theoretical logical complexity and throughput therefore represents a unique point where all descriptions of an emulated evolution are equivalent. Because of the equivalence of descriptions at this intersection point, it is not inconsistent to express different numbers of emulated carriages along a line of constant emulated carriage size.

$$E_{O_L} = \frac{R_i}{R_\mu} S = \frac{R_i}{R_\nu} \tag{7.2}$$

The analysis of 2-4-2 size systems reveals that we can characterize entire logical complexity/throughput space for any size system based on only two points - the minimum and maximum theoretical logical complexity/throughput combinations. Using these values, an evolution is describable in terms of the size of the emulated system with respect to the number of carriages and the average number of operational carriages. These characterizations are based on the slopes of the set of lines originating from the minimum theoretical logical complexity and throughput value that intersect the maximum theoretical throughput boundary and positions along these lines. These fundamental relationships describing the logical complexity/throughput space are illustrated qualitatively in Figure 7.50

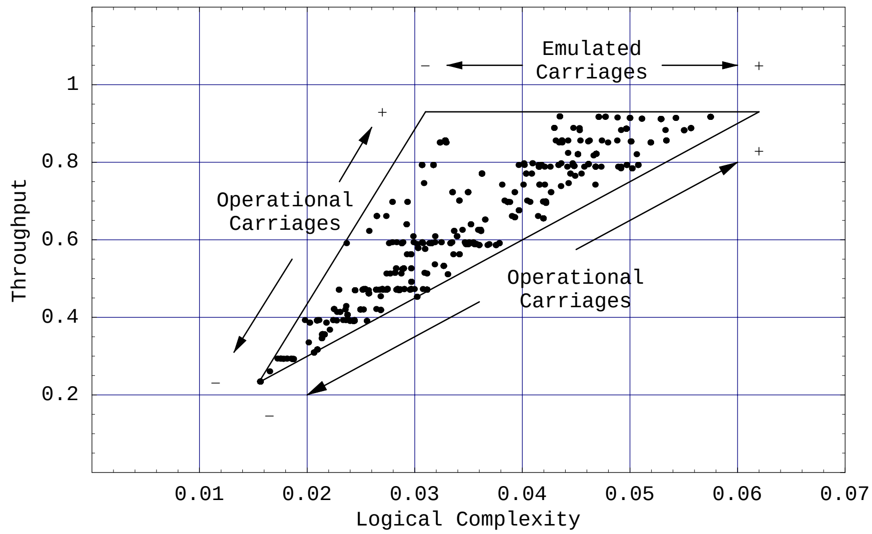

Figure 7.50: The characterization of the logical complexity/throughput space with respect to logical equivalence is based on the minimum and maximum theoretical logical complexity/throughput combinations. The slope of the line drawn from the minimum theoretical logical complexity/throughput combination through an evolution describes the number of emulated carriages and the distance along the line indicates the average number of operational carriages.

As defined, emulation describes the logical equivalence of an evolution and does not necessarily reflect the operational equivalence as a mimic might. However, the distinction between logical and operational equivalence is fundamental to the characterization of the logical complexity/throughput space. The qualitative characterization described by Figure 7.50 should in no way be interpreted to indicate the actual number of carriages and their average operational periods, but the number of carriages and their average operational periods *that result in a logically equivalent evolution.* A description of an evolution in terms of operational equivalence comes closer to reflecting the actual (average) values and requires a different characterization of the logical complexity/throughput space that is not interchangeable with characterization with respect to logical equivalence. In terms of characterization, operational equivalence is quite similar to the definition of mimics. However, operational equivalence is a broader term, as an evolution unique to a system size is expressible as an operational equivalent. To illustrate the distinction between operational and logical equivalence, consider an evolution with the lowest possible logical complexity at the maximum throughput (the 2-2-2 emulation example discussed previously). The single phase lag results in an emulation of a 2-2-2 evolution, but clearly the evolution is not operationally equivalent to a configuration with 2 carriages because all 4 carriages are utilized (for nearly the entire evolution), resulting in a throughput unachievable for a configuration with 2 carriages. Similarly, an evolution that emulates a 2-4-2 configuration with an average of half of the carriages operational throughout the evolution is operationally equivalent to a 2-2-2 configuration with both carriages operational and separated by a single phase lag. We might therefore also think of a characterization of the logical complexity/throughput space in terms of operational equivalence where lines of constant logical complexity indicate the number of phase lags present (as for logical equivalence at the maximum theoretical throughput) and lines of constant throughput indicate the average number of operational carriages. The equations describing operational equivalence parameters are based simply on the proportionality of an evolution's logical complexity and throughput to the corresponding minimum theoretical values. The parameters are equivalently determined from the maximum theoretical logical complexity and throughput. The relationships are presented in Equations 7.3 and 7.4.

$$E_{C_O} = \frac{R_i}{R_\nu} = S\frac{R_i}{R_\mu} = \frac{(S-1)R_i}{R_\mu - R_\nu} \tag{7.3}$$

$$E_{P_O} = \frac{(S-1)C_{x_i}}{C_{x_\mu} - C_{x_\nu}} - 1 \tag{7.4}$$

Qualitatively, this characterization is illustrated in Figure 7.51. A result of the ability to characterize the logical complexity/throughput space in different terms that might not always agree is that the actual number of carriages involved in an evolution and their average operational periods may be ambiguous. For example, the evolution with half the maximum theoretical throughput along the maximum logical complexity boundary could involve 4 carriages with various phase lags that halt at various points to result in an average operational period half of the maximum or it could involve 2 carriages operating throughout the entire evolution with a constant single phase lag. For cases like these, explicit analysis of the evolution, or at least its results, is required.

The characterization of the logical complexity/throughput space with respect to the definitions of Equations 7.1 and 7.2 are based on the determinism of the simulations. With non-deterministic cycle times, the concepts of minimum and maximum values are no longer valid with respect to logical complexity. A single carriage with four possible states that completes the delivery of a set of items with low cycle times can conceivably have a greater throughput and logical complexity than a system with four asynchronous operational carriages. If each individual cycle time has a finite distribution however, then we can define boundaries in the same manner as for a deterministic system, although the shape and number of the boundaries will change depending on the variability of the cycle times. For instance, the minimum throughput and logical complexity correspond to a minimum number of operational carriages with the combination of the greatest possible cycle times, but a maximum throughput with the minimum possible logical complexity represents a new point corresponding to the minimum number of operational carriages with the lowest combination of individual cycle times. Even though introducing variability into the cycle times results in a

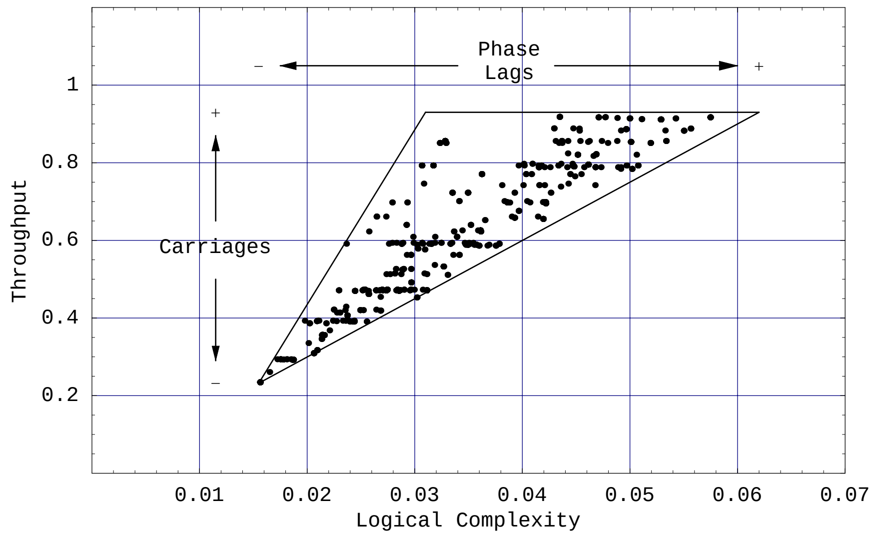

Figure 7.51: The qualitative characterization of the logical complexity/throughput space with respect to operational equivalence. Greater throughputs correspond to greater numbers of operational carriages and greater logical complexities correspond to more phase lags. The operational equivalent of an evolution does not necessarily reflect the actual dynamics of an evolution.

closer representation of reality, non-deterministic simulations are not analyzed in detail because non-determinism does not fundamentally alter the relationships between the definition of complexity and performance, characterization of which is essential to this work.

7.4.3 Algorithmic Complexity

Recognizing that the deterministic nature of the evolutions does not fundamentally change the characterization of the logical complexity/throughput space, we can apply the qualitative definition of algorithmic complexity to determine if the relationships are valid. We have seen that the slopes of the lines of constant carriage emulation are always positive and finite (a vertical line corresponds to independent carriages, which are always mimics and according to our definition, do not correspond to any emulated system). We therefore always have a positive relationship between logical complexity and throughput.

The validity of this relationship depends on the interpretation of the definition of algorithmic complexity. The strict definition of algorithmic complexity is based on the most compressed form of the evolution. On the surface, the addition of phase lags implies additional states to define for a complete description, which results in less possible compression and therefore higher algorithmic complexity. However, the equivalent carriage emulation and number of phase lags are *average* values for an evolution. Similarly, and with a greater impact on algorithmic complexity, the effects of averaging also arise along the lines of constant carriage emulation. In the application of the definition of logical complexity, the sequencing of the halting of carriages is ignored and the dynamics of evolutions are lost because logical complexity considers the evolution as a whole and is incapable of identifying patterns. An evolution that emulates a system with less than the maximum number of carriages, but has complex dynamics with many patterns will have a lower logical complexity than an evolution with a single repetitive pattern consisting of the maximum number of phase lags. By the strict definition of algorithmic complexity however, the complexity of the evolution emulating a smaller number of carriages is greater than the repetitive evolution because each pattern in the emulated system requires distinct information regarding the states involved and the sequence of the pattern in the evolution. So, in terms of algorithmic complexity, the actual states, no just their progression, appears relevant.

Since repetitive patterns result in the lowest algorithmic complexity, the evolution corresponding to the maximum theoretical logical complexity and throughput is only marginally more algorithmically complex than the evolution corresponding to the minimum theoretical logical complexity and throughput because the greater number of carriages involved in the evolution with maximum throughput implies more information for a complete description, while evolutions along a line of constant carriage emulation, which contain a mixture of evolutions and more diverse patterns have significantly greater algorithmic complexity. The distribution of evolutions with respect to the strict definition of algorithmic complexity is therefore fundamentally different than the distributions with respect to logical complexity, with all evolutions with simple repetitive patterns having essentially identical complexity, but ranging across all possible throughputs. Evolutions with halting carriages that contain mixtures of simpler evolutions and have multiple operational patterns have greater complexity, but at moderate levels of throughput. We might therefore construct a conceptual theoretical distribution of the algorithmic complexity/performance space, illustrated in Figure 7.52. We will refer to this conceptual distribution repeatedly when questioning the validity and practicality of various complexity measures, at the same time offering more description and justification of the shape of the boundaries.

The definition of algorithmic complexity implies that the maximum complexity in a deterministic system corresponds to an evolution where all carriages halt (except one, of course), resulting in the greatest number of distinct patterns requiring description. The maximum throughput for such a scenario corresponds to an evolution where carriages halt at the last possible unique moments. Specifically, where carriages halt at the $\tau - 1^{th}$, $\tau - 2^{th}$, ..., $\tau - (n-1)^{th}$ system cycle, where τ is the number of system cycles and n is the number of carriages. The throughput of any evolution with halting is always less than the maximum achievable for the evolutions in the set. In the search for optimal configurations (based on throughput and ignoring robustness for the moment), this effect implies the strict definition of algorithmic complexity results in a distribution of little value. A high algorithmic complexity indicates less than optimal throughput while low algorithmic complexity is ambiguous, since it corresponds to a range of simple patterns. For conventional systems, this effect also indicates that optimal performance with respect to throughput always corresponds to the

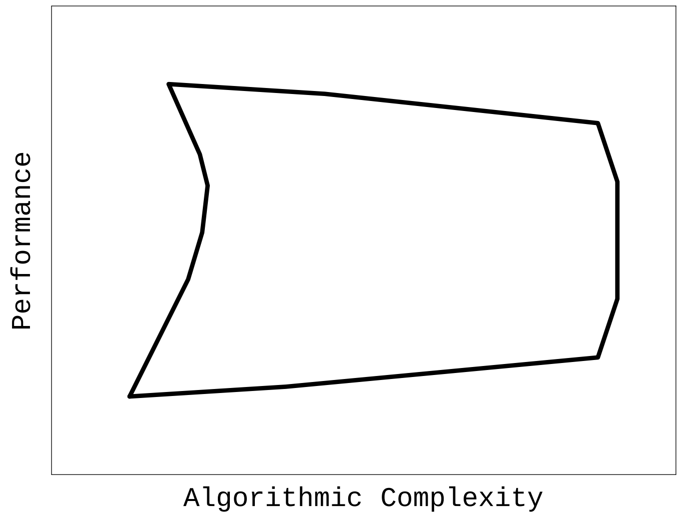

Figure 7.52: The conceptual distribution of algorithmic complexity and performance. Evolutions that have the greatest algorithmic complexity theoretically do not have maximum performance because they have greater exploration of their state spaces. The maximum theoretical throughput corresponds instead to simple evolutions with repetitive patterns that involve all carriages. The low algorithmic complexity of these evolutions with respect to system states is slightly, but definitely greater than evolutions with simple repetitive patterns of a single carriage. A strict interpretation of algorithmic complexity offers little practical value towards identifying optimal configurations with respect to performance.

simplest behavior.

By ignoring the patterns that arise in the course of an evolution and relaxing the definition of complexity to consider the information required to step through an evolution as with logical complexity, then optimality with respect to throughput is identifiable when equivalent carriage emulation is considered. So, in some respects, complexity with respect to average evolution characteristics provides a more useful format in the characterization of evolution distributions and in the search for optimal configurations. However, optimality is not strictly defined by throughput. Adaptability is essential in a complex system, and to see if the use of average evolution characteristics results in identification of optimal configurations with respect to adaptability, we have to investigate the various evolution subsets defined for 2-2-2 size systems.

Figure 7.53 presents the two-dimensional distributions with respect to logical complexity and throughput for the non-halting evolutions, complete evolutions, and the evolutions of robust configurations. The corresponding frequency landscapes are presented in Figure 7.54. The distributions of evolutions are fairly similar, although the non-halting distribution is slightly fuller than the other distributions, primarily because of the substantial difference in the number of evolutions in the set (1217 non-halting evolutions compared to 856 complete evolutions and 720 robust evolutions). The fuller distribution implies that many incomplete evolutions do not belong to high frequency clusters.

The frequency landscapes have similar shapes, with significant clustering at low to moderate values of logical complexity and throughput. Other peaks occur at the minimum logical complexity/throughput combination and at the minimum and maximum logical complexities at the maximum throughput. More detail regarding the frequencies is obtained from the histograms with respect to logical complexity and throughput, shown in Figure 7.55. The histograms reveal the clustering of logical complexities, which scales to the remaining frequencies fairly proportionally as the set is refined. The throughput histograms reveal a single dominant throughput, which drops disproportionally with set refinement. The throughput histograms also reveal fairly constant secondary peaks at throughput values of approximately 0.6, 0.7, 0.8, and 0.9, which implies the corresponding evolutions are common to all sets, although the corresponding logical complexities may be distributed.

Since most logical connectivities are clustered around a relatively low value and the frequencies drop approximately proportionally, the mean logical complexity increases as the set of evolutions is refined. Similarly, the mean throughput is greater for the robust configurations and complete evolutions because most incomplete and non-robust evolutions belong to the low throughput cluster. The mean logical complexities and throughputs for these sets are presented in Table 7.1.

Table 7.1: The mean logical complexity and throughput for the various 2-4-2 size evolution sets. Both mean logical complexity and throughput increase as evolutions become more robust, indicating a correlation between adaptability, complexity, and performance. Higher values for unique evolution subsets implies that increasing the system size results in more adaptable and more complex evolutions with greater potential performance.

	mean C_L	mean R
Non-halting	0.0335	0.5970
Complete	0.0342	0.6061
Robust	0.0345	0.6174
Unique Non-halting	0.0354	0.6447
Unique Complete	0.0361	0.6561
Unique Robust	0.0423	0.7336

The increase in logical complexity from the set of non-halting evolutions to the set of complete evolutions suggests a correlation between complexity and robustness with respect to the controllability of queues. Evolutions capable of handling a variety of input streams tend to have greater complexity. In addition, a greater mean throughput implies that the more robust evolutions are also better performers. The mean logical complexity and throughput for the evolutions corresponding to the most robust configurations indicate a similar correlation. *Configurations* that complete delivery for all queue distributions considered tend to have

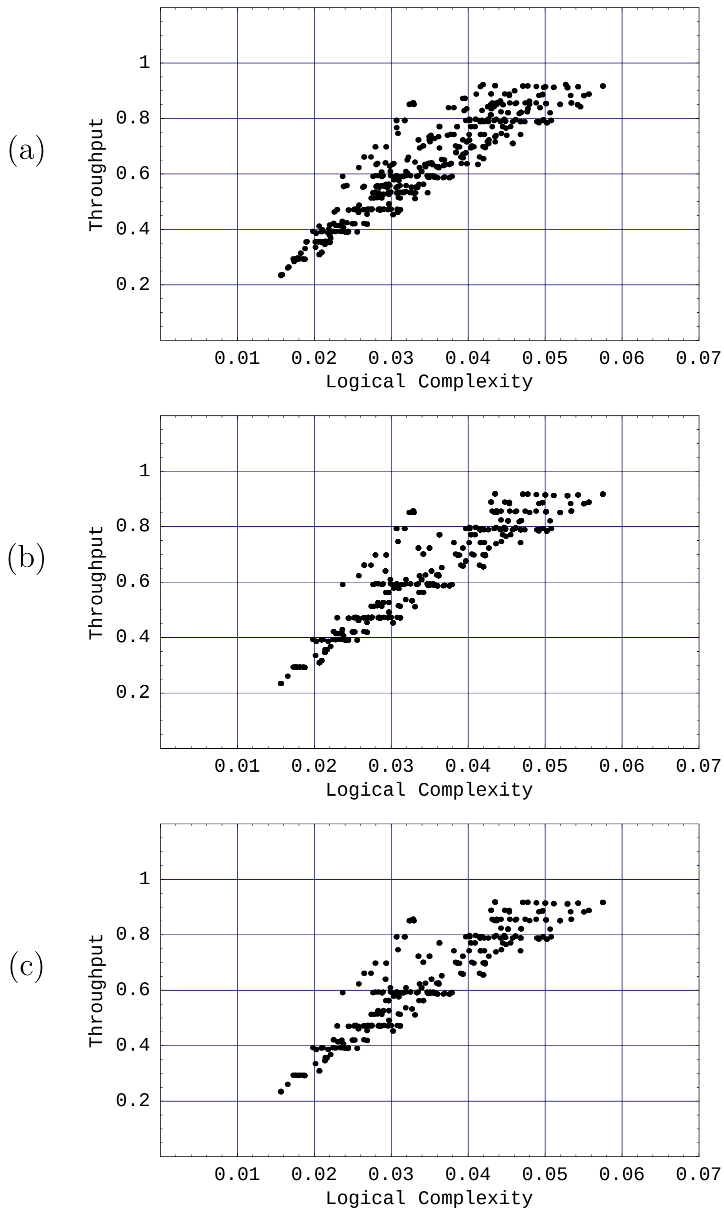

Figure 7.53: The two-dimensional distribution of 2-4-2 evolutions for (a) non-halting evolutions, (b) complete evolutions, and (c) the evolutions of the most robust configurations. The shapes of all three distributions are approximately identical, but the non-halting evolutions have more diverse locations, partly due to the larger number of evolutions in the set.

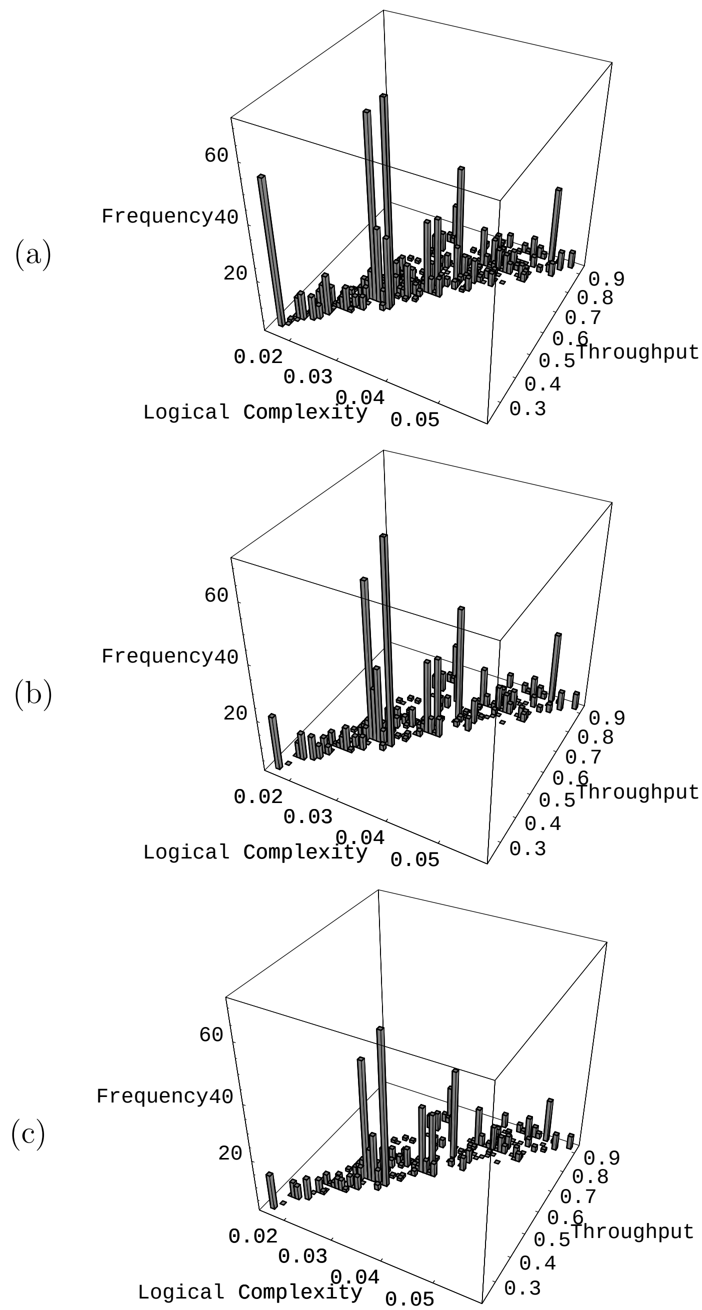

Figure 7.54: The three-dimensional frequency landscapes for (a) non-halting evolutions, (b) complete evolutions, and (c) the evolutions of the most robust 2-4-2 configurations. Clustering of evolutions at low to moderate values of logical complexity and throughput occurs in all distributions.

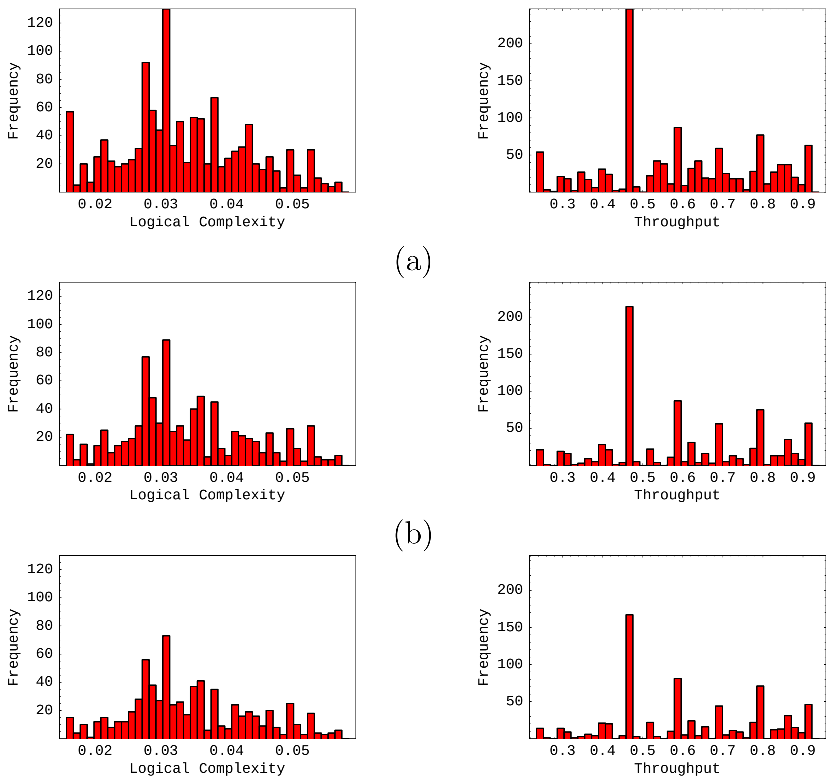

(c)

Figure 7.55: The histograms of non-unique 2-4-2 evolutions with respect to logical complexity and throughput for (a) non-halting evolutions, (b) complete evolutions, and (c) evolutions of robust configurations. The logical complexity distributions are fairly proportional, but the peak throughput drops consistently as the set is refined.

greater complexity and performance, but not necessarily the maximum values. As with 2-2-2 size systems, if we substitute the term adaptable for robust, then a comparison of the mean logical complexities and throughputs indicate that adaptability is directly related to both complexity and performance.

Table 7.1 also presents the mean logical complexity and throughput values for the evolution sets unique to 2-4-2 size systems. The qualitative differences between unique evolution sets are identical to the differences between the corresponding non-unique evolution sets, again implying a direct association between adaptability, complexity and performance. The greater absolute values of logical complexity and throughput for any unique evolution set compared to any non-unique evolution set also suggest a relationship between these values and system size. Larger systems are able to support greater complexity and, by the observed relationship with throughput, will also have greater performance on average. Greater complexity also suggests that larger systems also have greater adaptability. This relationship is especially true in configurations where the number of shafts is greater than both the number of queues and magazines - for given numbers of queues and magazines, more carriages result in a greater probability of collateral circuits.

The two-dimensional distributions for the unique evolution sets are presented in Figure 7.56 and show that the distributions for unique non-halting and complete evolution sets have a similar shape to the non-unique evolution distributions. However, the distribution of unique robust evolutions is only a fraction of the other distributions. This is partly because of the relatively small number of robust configurations. Only 60 evolutions comprise the set of unique, robust evolutions compared to 881 unique, non-halting evolutions and 595 unique, complete evolutions. While the distribution of unique, robust evolutions provides an explanation of the relatively high mean logical complexity and throughput, three-dimensional frequency landscapes and histograms are required for unique non-halting and complete evolution sets to characterize their salient features. The three-dimensional frequency landscapes for the unique evolution sets are presented in Figure 7.57 and the corresponding histograms of the logical complexity and throughput cross sections are presented in Figure 7.58.

The principal differences evident between these histograms and the histograms for the non-unique evolution sets, beyond the small number of evolutions comprising the evolutions of robust configurations in Figure 7.58(c), occur at the peak logical complexities and throughput of non-unique evolution sets. The lower frequencies corresponding to unique evolution sets are not necessarily attributed to a reduction in the number of evolutions comprising a set. There are 881 evolutions in the unique, non-halting set compared to 856 non-unique, complete evolutions. The largest cluster of non-unique, complete evolutions contains 89 evolutions between logical complexities of 0.0303 and 0.0313. Despite the comparable total number of evolutions in the set, there are only 30 unique, non-halting evolutions in the same range of logical complexities. Similarly, 219 non-unique, complete evolutions have throughputs between 0.458 and 0.475 while there are only 80 unique, non-halting evolutions in this range.

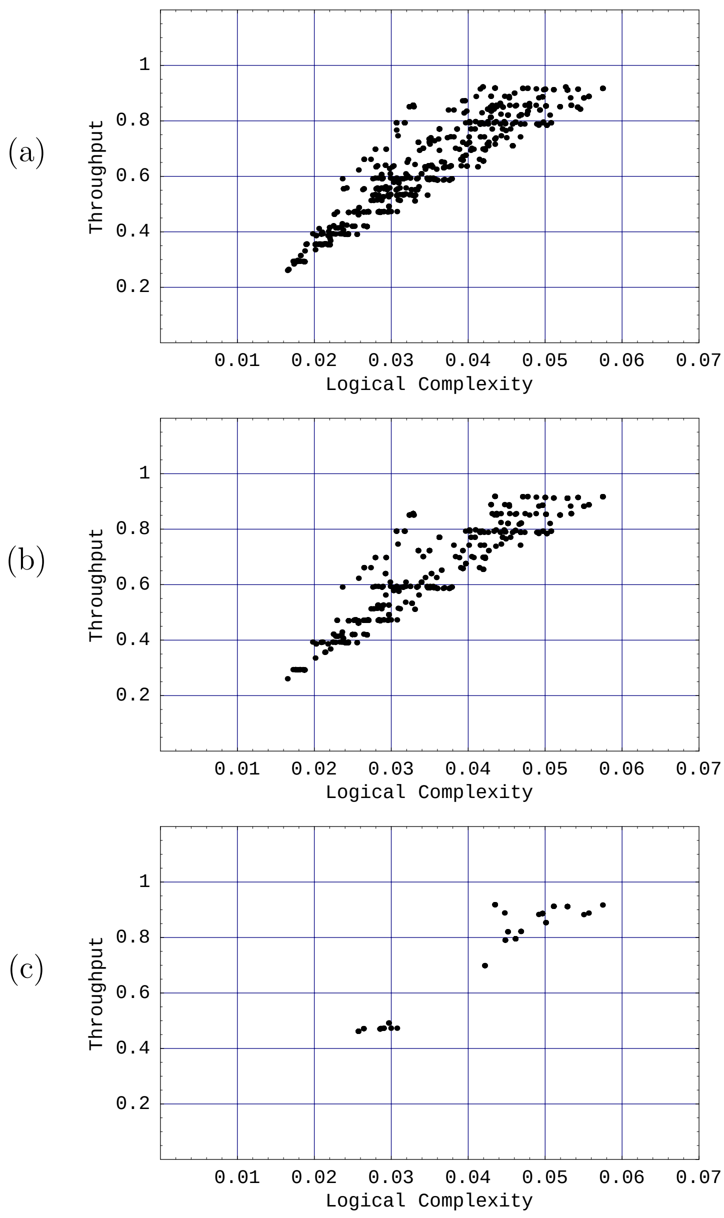

Figure 7.56: The two-dimensional distribution of unique 2-4-2 evolutions for (a) non-halting evolutions, (b) complete evolutions, and (c) the evolutions of the most robust configurations. The non-halting and complete evolution distributions occupy a similar space as the non-unique distributions, but the unique robust evolutions exist only at higher logical complexities and throughputs.

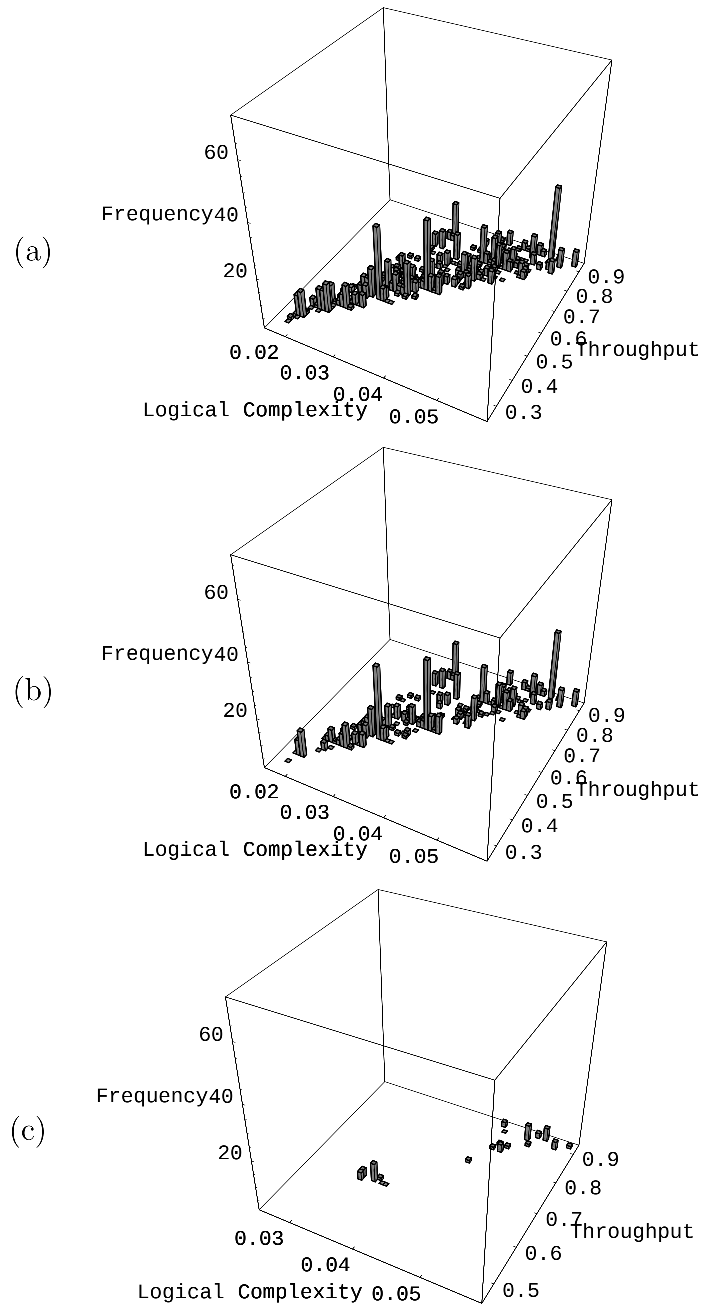

Figure 7.57: The three-dimensional frequency landscapes for (a) non-halting evolutions, (b) complete evolutions, and (c) the evolutions of the most robust 2-4-2 configurations. The frequency landscapes help explain the increased mean logical complexity and throughput found in unique evolution sets that is not apparent in two-dimensional distributions.

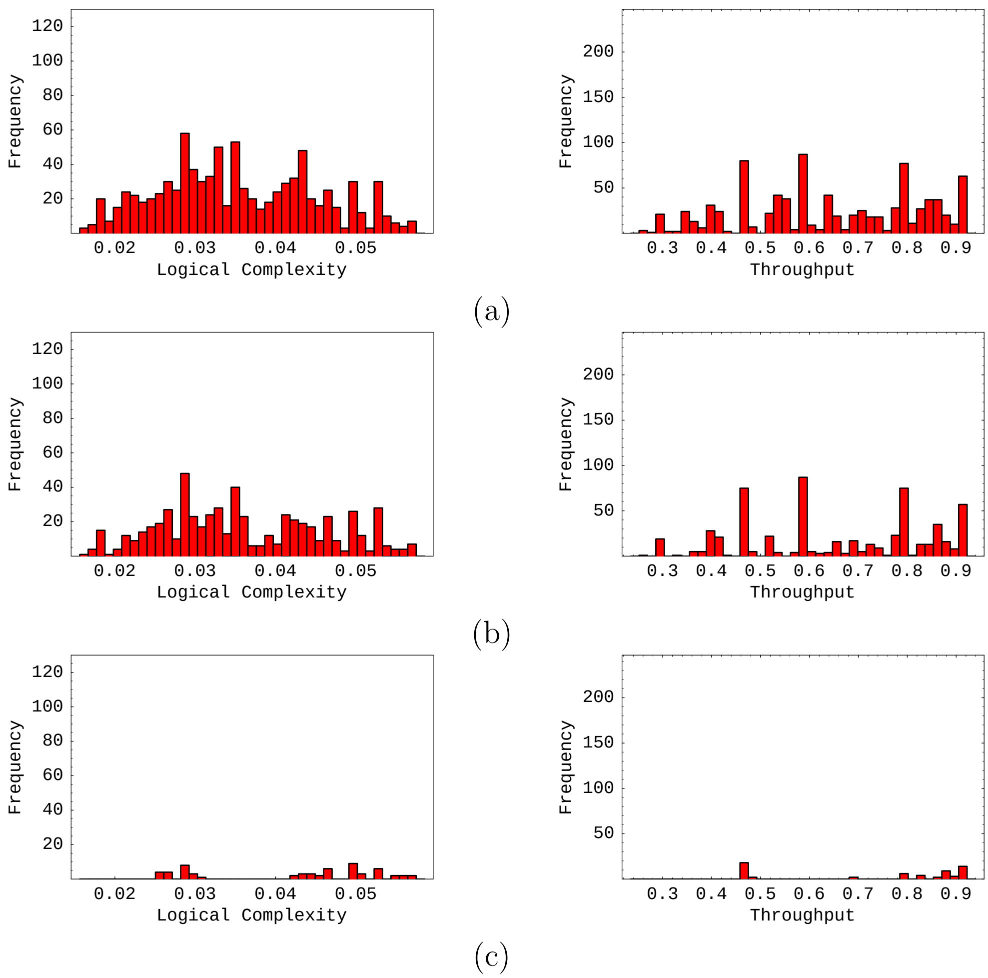

Figure 7.58: The histograms of unique 2-4-2 evolutions with respect to logical complexity and throughput for (a) non-halting evolutions, (b) complete evolutions, and (c) evolutions of robust configurations. The evolution clusters at the low to moderate logical complexity and throughput found in non-unique evolution sets are significantly reduced in unique evolution sets as is the throughput peak at a value of 0.7. These missing evolutions correspond to mimics of 2-2-2 and 2-3-2 systems.

The characteristics of the mimics that constitute the high frequency clusters can be described in terms of operational equivalence defined in Equations 7.3 and 7.4. They could also be described in terms of logical equivalence, although operational equivalence provides a closer representation of reality. However, ambiguity with respect to the actual evolution is always present regardless of the equivalent description. A large fraction of the mimics near the high frequency cluster have a logical complexity of 0.0311174 and a throughput of 0.471476. By Equations 7.3 and 7.4, these evolutions are operationally equivalent to an evolution with 2.03 operational carriages with 1.01 phase lags. Given the near integer equivalent values, it is safe to assume these evolutions mimic a 2 carriage configuration, with phase-locked, phase-lagged carriages, which is identical to the maximum logical complexity/maximum throughput evolution for 2-2-2 systems. A look at the compressed carriage state histories viewed collectively in Figure 7.59 and individually in Figure 7.60 for one of these evolutions (65387 (0-100)) verifies the accuracy of the operational equivalence. Only two carriages are in operation and are separated by a single phase lag, the maximum possible for two operational carriages.

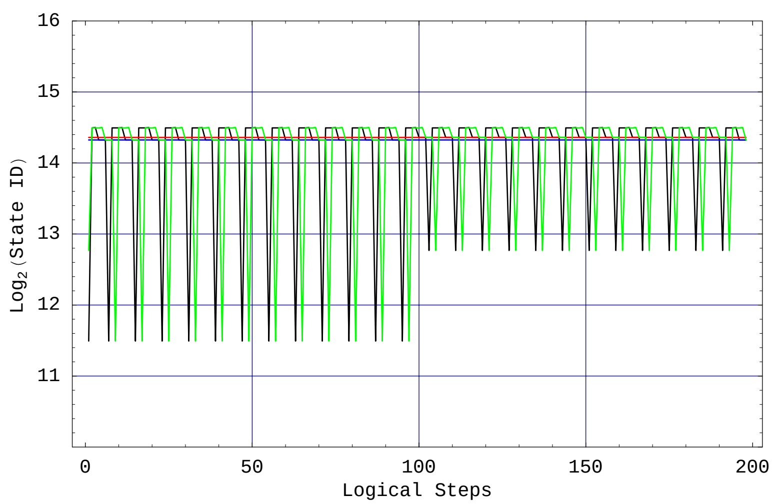

Figure 7.59: The collective compressed carriage state histories for evolution 65387 2-4-2 (0-100). Only a single phase lag exists in the actual evolution which matches the description with respect to operational equivalence.

The characteristics of other mimics at other throughputs and logical connectivities are also expressible in terms of operational equivalence. For instance, the cluster of evolutions at throughputs between 0.682 and 0.699 correspond to an operational equivalent of 3 carriages. Evolutions at these throughputs are therefore most likely mimics of 2-3-2 evolutions.

As with the 2-2-2 configurations, the unique robust 2-4-2 configurations share the trait of complete connectivity between shafts and magazines. Regardless of the connectivity between shafts and queues, all items will always be delivered, although shaft utilizations may be unequal as carriages halt because of limited access to queues. Unlike 2-2-2 systems, the

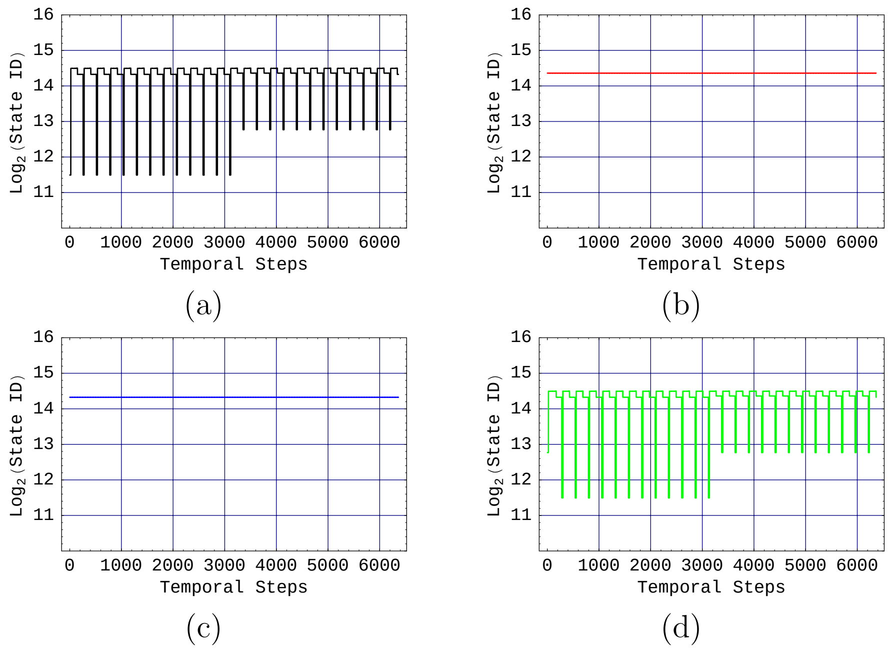

Figure 7.60: The individual temporal evolution histories for (a) carriage 1, (b) carriage 2, (c) carriage 3, and (d) carriage 4 for configuration 65387 2-4-2 with a (0-100) queue distribution. Two of the carriages are operational for the entire evolution while the other two are never utilized because they are only connected to the first magazine. The description of the number of carriages with respect to operational equivalence indicates that only two carriages (on average) are operational.

unique robust configuration with complete connectivity (65535) does not correspond to the maximum logical complexity or throughput for the set. With greater connectivity and more carriages comes the chance for more carriage interactions which tend to increase the length of an evolution because of more delays without additional logical complexity from more state transitions. Once again, it is important to define what is meant by optimality in an optimal search. While these unique robust configurations are not necessarily optimal with respect to throughput, they are adaptable in the sense that they are guaranteed to complete delivery regardless of the input stream. Similarly, non-unique robust configurations are optimal with respect to adaptability, although they may mimic smaller systems with even lower performance to achieve complete delivery.

7.5 Other Larger Systems

The relationships between logical complexity and throughput for evolution sets of other system sizes are fundamentally similar to those discussed for 2-2-2 and 2-4-2 systems. The minimum theoretical logical complexity/throughput combination always corresponds to a single carriage transporting all items from queues containing only one item type. The maximum theoretical logical complexity/throughput combination always corresponds to all carriages operating with the greatest number of phase lags. The extent to which the maximum theoretical values are attained is dependent on the distribution of item types in the queues and the relative numbers of queues, shafts, and magazines. The relative numbers of queues, shafts, and magazines also affects the shape of the distribution of evolutions.

Characterization of the logical complexity/throughput space with respect to logical and operational equivalence is possible for all system sizes, based on the minimum and maximum theoretical logical complexities and throughputs, although the absolute values of the slopes of lines of constant emulated carriage size (for a given nominal carriage size) may be different for different system sizes. However, the actual slope of the maximum logical complexity boundary remains fairly constant for all systems, assuming deterministic cycle times.

Characterization of 2-2-2 and 2-4-2 systems revealed correlations between adaptability, logical complexity, and throughput. These correlations are based on a comparison of the mean values of logical complexity and throughput for evolution subsets. Tables 7.2 and 7.3 present the mean logical complexities and throughputs for the evolution subsets for all system sizes considered and indicate the same trends seen in 2-2-2 and 2-4-2 systems. The mean logical complexity and throughput for complete evolutions are always greater than the corresponding means for non-halting evolutions, implying a correlation between robustness, logical complexity, and throughput, since complete evolutions are arguably more robust than incomplete evolutions in the context of uncontrollable queues. Similarly, evolutions corresponding to the most robust configurations always have greater mean logical complexity and throughput than for complete evolutions, which also implies a correlation between robustness, logical complexity, and throughput. Adaptability with respect to tolerance to variety of queue distributions corresponds to greater logical complexity and performance. However, the most adaptable configurations do not necessarily have the maximum logical complexity and throughput, indicating a trade-off between performance and adaptability and raises the question of what is optimal in a complex system. Whereas reductionist analyses search for optimality with respect to some performance measure, in complex systems, adaptability arguably equates to optimality.

The mean logical complexities and throughputs for evolution sets consisting of unique evolutions follow the same trends as the corresponding non-unique sets, supporting the conclusions reached regarding the correlations between adaptability, logical complexity, and throughput. The mean values of the unique sets are always greater than the corresponding non-unique sets and, in many cases, the mean values for unique, non-halting evolutions are greater than the mean values for the non-unique, robust evolutions. The greater values indicate that the smaller systems represented by mimics have lower complexity and performance, suggesting that larger systems are able to support greater logical complexity and throughput. If logical complexity and throughput are related to adaptability, then greater logical complexity and throughput in larger systems imply that larger systems are also more adaptable.

Table 7.2: The mean logical complexity for evolution subsets of different system sizes (N = non-halting, C = complete, R = robust, UN = unique and non-halting, UC = unique and complete, UR = unique and robust, M = mimics). For all system sizes, the mean logical complexity for complete evolutions is greater than for non-halting evolutions and the mean logical complexity for robust evolutions is greater than for complete evolutions. The same trend occurs for unique evolution sets. In most cases, the mean for unique, non-halting evolutions is greater than the mean for non-unique, robust evolutions. The increases in the mean indicate a correlation between adaptability and logical complexity. Greater mean logical complexities for unique evolution sets also imply that larger systems support greater complexity and are therefore more adaptable.

System	N	C	R	UN	UC	UR	M
1-2-2	0.02574	0.02574	0.02574	0.02773	0.02773	0.03081	0.01577
1-2-3	0.02491	0.02491	0.02491	0.02695	0.02695	0.03081	0.01577
1-2-4	0.02475	0.02475	0.02475	0.02671	0.02671	0.03081	0.01577
2-2-2	0.02347	0.02639	0.02673	0.02519	0.02829	0.03112	0.01570
2-2-3	0.02332	0.02492	0.02531	0.02388	0.02701	0.03112	0.01570
2-3-2	0.02856	0.02917	0.02921	0.02981	0.03017	0.03231	0.02457
2-3-3	0.02742	0.02841	0.02884	0.02951	0.02994	0.03261	0.02285
2-3-4	0.02718	0.02820	0.02893	0.02943	0.02992	0.03276	0.02250
2-4-2	0.03357	0.03416	0.03446	0.03541	0.03614	0.04216	0.02874
2-4-3	0.03249	0.03315	0.03391	0.03540	0.03594	0.04291	0.02812
3-2-2	0.02170	0.02448	0.02492	0.02306	0.02581	0.02700	0.01565
3-3-2	0.02775	0.02968	0.02983	0.02927	0.03174	0.03657	0.02306
3-4-2	0.03194	0.03309	0.03293	0.03317	0.03427	0.04016	0.02878
4-2-2	0.02105	0.02388	0.02457	0.02229	0.02498	0.02602	0.01562
4-2-3	0.02041	0.02288	0.02366	0.02159	0.02431	0.02602	0.01562
4-2-4	0.02027	0.02234	0.02304	0.02136	0.02385	0.02602	0.01562
4-3-2	0.02691	0.02915	0.02939	0.02828	0.03070	0.03362	0.02274
4-3-3	0.02562	0.02799	0.02860	0.02764	0.02985	0.03363	0.02131
4-4-2	0.03179	0.03427	0.03419	0.03331	0.03644	0.04496	0.02781

Table 7.3: The mean throughput for evolution subsets of different system sizes (N = non-halting, C = complete, R = robust, UN = unique and non-halting, UC = unique and complete, UR = unique and robust, M = mimics). The same trends with respect to increases in mean values with set refinement evident for logical complexity apply to the mean throughput. Movement of the logical complexity and throughput in the same direction implies a correlation - more logically complex evolutions also have greater throughput. Additionally, if adaptability is related to logical complexity, then the most adaptable configurations have, on average, better performance.

System	N	C	R	UN	UC	UR	M
1-2-2	0.3860	0.3860	0.3860	0.4159	0.4159	0.4621	0.2365
1-2-3	0.3736	0.3736	0.3736	0.4042	0.4042	0.4621	0.2365
1-2-4	0.3713	0.3713	0.3713	0.4006	0.4006	0.4621	0.2365
2-2-2	0.3748	0.4015	0.4069	0.4058	0.4310	0.4715	0.2355
2-2-3	0.3616	0.3800	0.3859	0.3913	0.4126	0.4715	0.2355
2-3-2	0.4849	0.4884	0.4915	0.5208	0.5195	0.5549	0.3704
2-3-3	0.4670	0.4713	0.4832	0.5172	0.5117	0.5464	0.3574
2-3-4	0.4653	0.4673	0.4852	0.5174	0.5102	0.5411	0.3563
2-4-2	0.5970	0.6061	0.6174	0.6447	0.6561	0.7336	0.4722
2-4-3	0.5780	0.5873	0.6091	0.6509	0.6583	0.7518	0.4686
3-2-2	0.3492	0.3710	0.3780	0.3751	0.3916	0.4078	0.2348
3-3-2	0.4796	0.5036	0.5102	0.5224	0.5550	0.6361	0.3479
3-4-2	0.5875	0.5992	0.6028	0.6300	0.6381	0.7065	0.4777
4-2-2	0.3407	0.3618	0.3713	0.3651	0.3782	0.3923	0.2343
4-2-3	0.3351	0.3468	0.3582	0.3599	0.3691	0.3923	0.2343
4-2-4	0.3347	0.3394	0.3492	0.3584	0.3630	0.3923	0.2343
4-3-2	0.4653	0.4923	0.5016	0.5054	0.5344	0.5810	0.3426
4-3-3	0.4456	0.4672	0.4827	0.4956	0.5151	0.5812	0.3395
4-4-2	0.5884	0.6188	0.6250	0.6381	0.6767	0.7910	0.4587