Chapter 10

Number of Unique States Used

10.1 Algorithmic Complexity

In the discussion of state complexity, we saw that normalization by the temporal evolution length introduced a correlation between state complexity and throughput and resulted in a divergence from a strict interpretation of algorithmic complexity.

At the minimum and maximum throughputs, the number of operational carriages is known. However, intermediate throughput values do not explicitly indicate the number of operational carriages and therefore the sequence of carriage haltings is unknown. The maximum potential number of state combinations is therefore assumed for all intermediate throughputs. The throughput does indicate some information regarding the *timing* of carriage haltings, which has a direct influence on temporal evolution length. Since both the throughput and state complexity are both inversely proportional to the temporal evolution length, a linear relationship between state complexity and throughput is inherent.

According to the definition of algorithmic complexity, the information required to describe the number of repetitions of states/patterns is negligible with respect to the information required to describe the states comprising system patterns, the number of distinct patterns, and their global sequencing. Normalization by the temporal evolution length therefore results in a divergence from a strict definition of algorithmic complexity and implies that description of complexity solely by the number of states is a closer representation of algorithmic complexity.

However, under the assumption of a constant minimum/maximum theoretical number of states, removal of normalization, while bringing complexity conceptually closer to algorithmic complexity, removes any correlation between complexity and throughput. The theoretical distribution is presented in Figure 10.1(a) for intermediate throughputs under this assumption and indicates that the information required to describe an evolution at a lower throughput can theoretically equal the information required to describe an evolution at a higher throughput, if both evolutions experience the same states, regardless of the timing, assuming the lower throughput evolution can enter the equivalent states as the higher throughput evolution before halting of carriages begins.

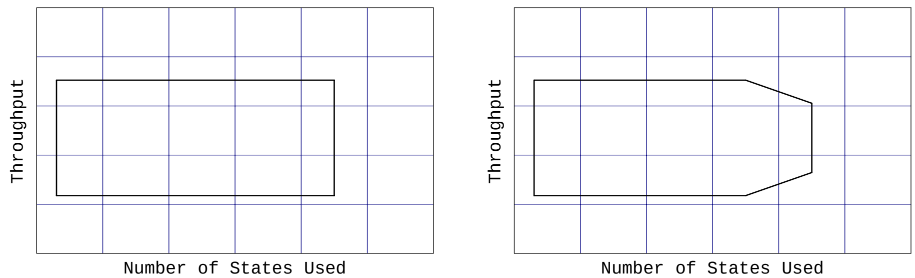

(a) (b)

Figure 10.1: Qualitative theoretical distributions of the number of unique states with respect to throughput. The minimum and maximum number of states boundaries in (a) are vertical, suggesting no correlation between the number of states and throughput exists because, across the range of intermediate throughput values, evolutions can theoretically enter the same number of states, but carriages may halt at different times. The boundaries can be refined in (b) at near minimum and maximum throughputs, because more information about the possible halting sequences is known.

While the theoretical distributions offer no apparent correlations between the number of states and throughput, actual distributions, like the example distribution of 2-3-4 size evolutions in Figure 10.2 do exhibit correlations. The reasons are related to several factors, including the accuracy of the theoretical boundaries and the validity of the assumptions used in their creation. More importantly, the size of the potential state space and the susceptibility to halting of carriages (and therefore throughput) are directly related to the connectivity of a configuration in the context of the queue distribution used.

While we know carriage halting must occur at all intermediate throughput values but the throughput values do not specifically indicate when halting occurs in an evolution, at lower and higher throughputs, we do have more information regarding halting, which refine the theoretical boundaries. At sufficiently low throughputs, we can calculate from the throughput that, given integer values of carriage cycles, only a certain fraction of the total number of carriages can be operational in addition to the over-utilized carriage. In most cases, it may be ambiguous which carriages are operational and for how many cycles. Regardless, the potential state space is less than that for an evolution where the inclusion of all carriages and subsets of operational carriages are possible - which can occur at intermediate throughputs. The same limitations apply at sufficiently high throughputs, where the throughput can indicate that only a certain fraction of the total number of carriages can be non-operational. With this additional information, the theoretical boundaries of the maximum number of states are refined, depending on the system size, at the minimum and maximum through-

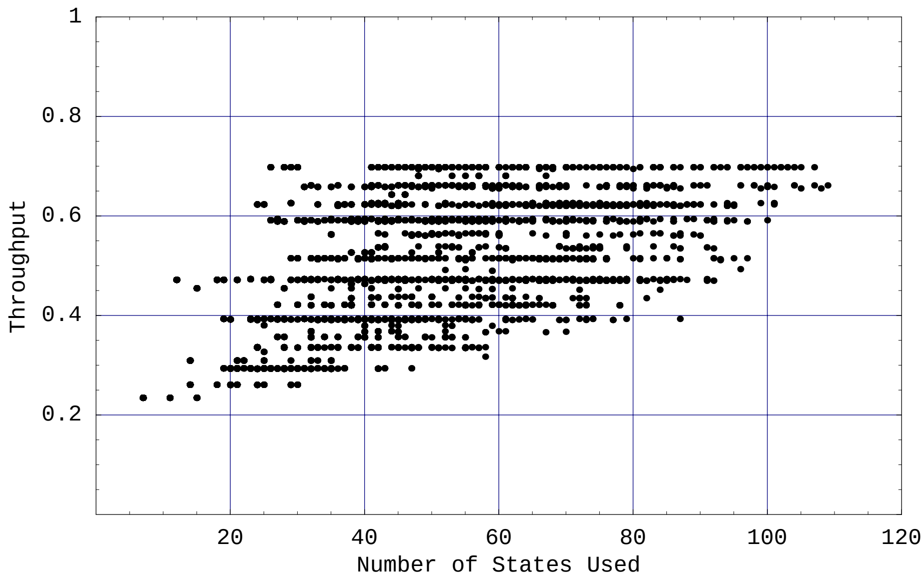

Figure 10.2: The distribution of complete evolutions of 2-3-4 size systems reveals a relationship between the number of unique states and throughput. A relationship exists despite the fact that the number of states closely follows the definition of algorithmic complexity and algorithmic complexity does not exhibit any theoretical relationship to performance.

puts as in Figure 10.1(b) to reflect the fact that not all combinations of carriage states are possible.

In the discussion of state complexity, we saw that the maximum theoretical state complexity significantly overestimates the actual maximum values at any throughput. The overestimation is a result of the ambiguous relationship between the timing of carriage haltings and throughput. Regardless of the effect of refinements on the theoretical boundaries due to additional information at sufficiently small and large throughputs, the theoretical boundaries offer few clues to explaining the actual relationships between the number of states and throughput evident in distributions such as that in Figure 10.2, simply because the minimum and maximum boundaries of the number of distinct system states assume all possible state combinations for a given throughput. The set of all possible state combinations is a significant overestimation of the actual number of states because it assumes non-determinism in evolutions and therefore includes unobtainable system states.

10.2 Qualitative Characterizations

In actual distributions, the non-trivial relationship between the number of states and throughput is explained by examining the conditions that result in halting carriages. Carriages halt as a result of a lack of deliverable items, which is a function of the connectivity of the configuration, in particular, the connectivity between shafts and magazines, described in the context of the queue distribution, described in terms of not only actual distributions, but also the amount of variety of item types. The relationship between variety in the queue distribution, connectivity, the number of states, and throughput is evident when we examine specific evolutions using the example 2-3-4 system size distribution presented in Figure 10.3.

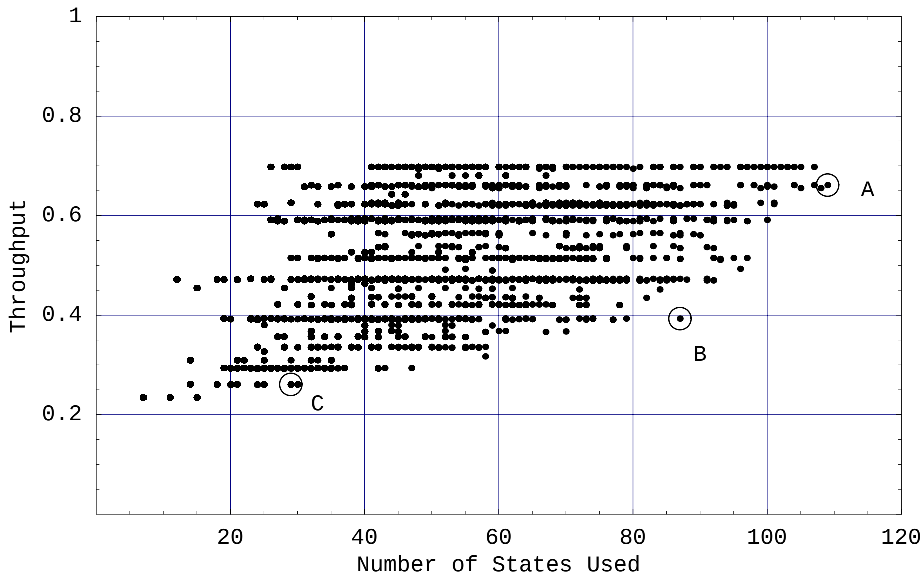

Figure 10.3: Points along the maximum number of unique states boundary illustrate the relationships between connectivity and variety in the queue distribution and establishes a correlation between the number of unique states and throughput. Points *(A)* (configuration 227261), *(B)* (configuration 251787), and *(C)* (configuration 94495) all correspond to evolutions with maximum variety in the queue distribution.

The evolution corresponding to point A is configuration 227261 with a queue distribution of 20-40-20-20. The throughput of this evolution is slightly less than the maximum because the second carriage halts two cycles prior to the end of the evolution and the third carriage halts one cycle prior to the end of the evolution, which is illustrated in the individual carriage histories for this evolution in Figure 10.4. Since halting occurs (not simultaneously), the system has the opportunity to enter more potential states, further illustrating the maximum complexity at sub-maximal throughput previously seen with respect to state and compressed state complexity. However, the more important factor in the exploration of the state space is the variety of item types in the queue distribution. Item types corresponding to all magazines are present, resulting in the potential for carriages to enter all magazines (and different queues depending on the connectivity) and forming more combinations of carriage states than if all items went to a single magazine. But variety in the queue distribution does not guarantee maximum exploration of the potential state space - variety must be taken in the context of connectivity. Additionally, the absolute distribution of item types is important, even if variety is constant, again in the context of connectivity.

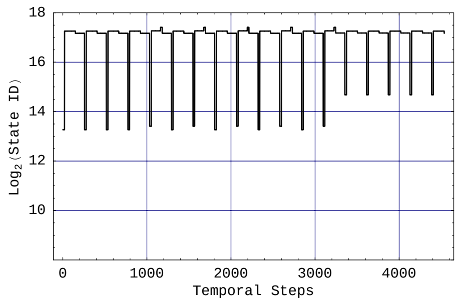

(a)

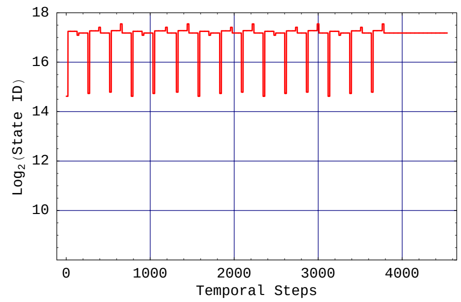

(b)

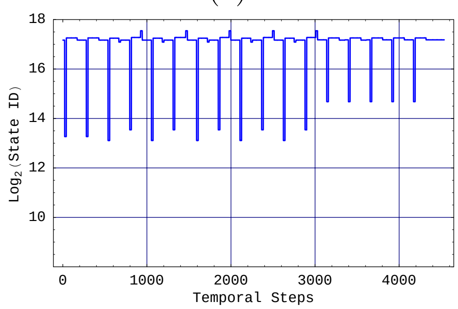

(c)

Figure 10.4: The individual carriage state histories for configuration 227261 with a 20-40-20-20 queue distribution, which corresponds to point A in Figure 10.3. Since halting occurs at the end of the evolution, but not simultaneously, the evolution enters the maximum number of states for this system size and has near maximum throughput.

The connectivity of configuration 227261 is presented in Figure 10.5 and reveals that the configuration is fairly well connected, with each shaft connected to three magazines (and each magazine connected to at least two out of three shafts), providing multiple paths between several queues and magazines. When the queue distribution is most uniform (20% of three item types and 40% of a single item type), the number of states is typically near maximal. Only for one of the most uniformly distributed queue distributions (40-20-20-20) is the number of distinct states relatively low. This queue distribution corresponds to an evolution with maximum throughput, where all carriages remain operational throughout the evolution.

$$(SQ) = \begin{pmatrix} 1 & 1 \\ 0 & 1 \\ 1 & 1 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 1 & 2 & 1 & 2 \\ 2 & 2 & 2 & 3 \end{pmatrix}$$

Figure 10.5: The incidence matrices for configuration 227261 2-3-4, which show that the first shaft is not connected to the first magazine. When the queue distribution is (100-0-0-0), the first shaft is not utilized despite complete connectivity to all queues, and the evolution mimics a two shaft configuration.

For the configuration corresponding to point B in Figure 10.3, configuration 251787 with a 20-40-20-20 queue distribution, fewer states are possible for a queue distribution with maximum variety because of a decrease in connectivity. Figure 10.6, which presents the connectivity for configuration 251787, reveals that only one shaft is connected to the second magazine and the second shaft is connected to a single magazine. For the 20-40-20-20 queue distribution, only the first carriage is capable of delivering the most abundant item type corresponding to the second magazine. The first carriage is also the only carriage connected to the third and fourth magazines that is also connected to the first queue. The first carriage must therefore transport all items bound for the second magazine from both queues and items corresponding to the third and fourth magazines from the first queue. This connectivity limits the state space entered because, while the first carriage enters all possible states except those corresponding to the first magazine (including loading items bound for the first magazine in either queue and movement towards the first magazine), the second carriage only enters states related to the first magazine and the third carriage never enters the second queue. At the same time, the limited connectivity also results in the halting of carriages two and three since they can only carry certain items loaded in particular locations. The connectivity, for this particular queue distribution, therefore results in less diversity in system states and lower throughput, which establishes a correlation between the number of states and performance.

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

Figure 10.6: The incidence matrices for configuration 251787 2-3-4. The first carriage is the only carriage that can transport items to the second magazine from both queues and that can transport items from the first queue to the third and fourth magazines. The second carriage only transports items to the first magazine and the third carriage only takes items from the second queue. The connectivity limits the number of potential system states for an evolution and also results in carriage haltings, although the exact limits are dependent on the actual items in the queue distribution. Through the connectivity, a relationship between complexity and throughput therefore exists.

To extend this concept of the association between sparse connectivity, limitation on the number of system states, and the early onset of carriage haltings to its logical limit, we can imagine a configuration that is capable of complete delivery of all queue distributions (with present connectivity between every queue and magazine), and either minimal valid connectivity in the SM or SQ matrix, within the definition of valid configuration. For the 2-3-4 system size, an example configuration is 94495, which has connectivity shown in Figure 10.7.

The first and second shafts/carriages share identical connectivity and are only able to transport items from the second queue to the fourth magazine. The third carriage must therefore transport all items bound for the first, second, and third magazines from both magazines

$$(SQ) = \begin{pmatrix} 0 & 1 \\ 0 & 1 \\ 1 & 1 \end{pmatrix} \quad (SM) = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \quad (QM) = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 3 \end{pmatrix}$$

Figure 10.7: The incidence matrices for configuration 94495 2-3-4 have minimal valid connectivity between shafts and magazines. While the evolution of such a system with any queue distribution will be complete, the potential state space is near minimal and carriages will not halt only for very specific queue distributions. This sparsely connected system further illustrates the link between the number of states, throughput, and connectivity.

and items bound for the fourth magazine in the first queue and is also capable sharing the load of the carriages with limited capabilities.

When the greatest diversity of item types is present in the queues, the first and second carriages still only experience four distinct (and identical) states each. The majority of system states therefore result from the evolution of the third carriage. The dominance of the third carriage states is evident from Figure 10.8, which presents the individual temporal carriage histories for the evolution of configuration 94495 with a (40-20-20-20) queue distribution, the evolution corresponding to point C in Figure 10.3. The first and second carriages halt relatively early in the evolution, after they exhaust all items bound for the fourth magazine in the second queue. The remaining number of *system* states are then a function of the variety of individual *carriage* states entered by the third carriage, which is dependent on the variety of items destined for the first, second, and third magazines (and items bound for the fourth magazine located in the first queue).

The only difference between evolutions of this sparsely connected configuration when the greatest diversity is present in the queue distribution occurs when there are sufficient items corresponding to the fourth magazine. In this scenario, the first and second carriages remain operational long enough for the third carriage to experience multiple patterns (due to de facto priority logic), resulting in more system states encountered than if the first and second carriages halted sufficiently early because they had fewer items to transport so that the first, second, and third carriages collectively create a single system pattern.

The example evolutions demonstrate that connectivity has a direct effect on the number of unique states entered by a system and when carriages halt in an evolution, and therefore introduces a direct relationship between the number of states and throughput. In general, we can say that configurations with greater connectivity tend to correspond to evolutions with a greater number of unique states and throughput because greater connectivity offers carriages more flexibility in item selection, which results in more uniform carriage utilization and lower throughput. The analogy in complex systems terminology is that greater connectivity provides greater adaptability for carriages to 'survive' throughout the course of an evolution. A sparsely connected configuration is composed of carriages that operate in niche environments, which 'die out' in the wrong environmental conditions (queue distributions).

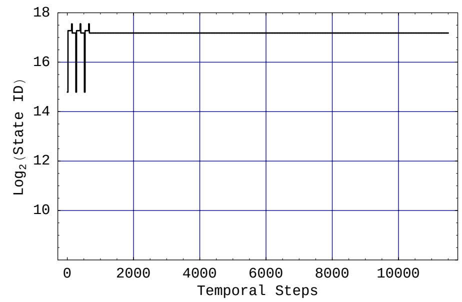

(a)

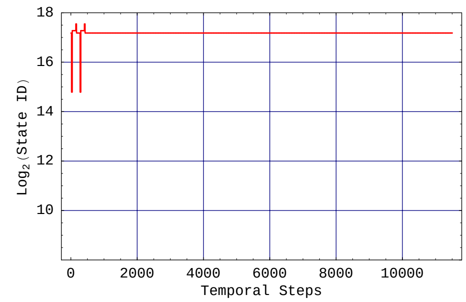

(b)

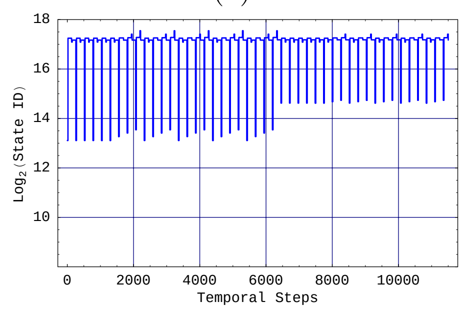

(c)

Figure 10.8: The individual carriage state histories for configuration 94495 with a 40-20-20-20 queue distribution, which corresponds to point C in Figure 10.3. Despite a queue distribution with the greatest amount of item variety, the sparse connectivity results in the first and second carriages halting early in the evolution, limiting the number of unique states and the throughput.

However, placing the connectivity in the context of the 'wrong' environmental conditions is key to the relationship between connectivity, the number of states, and throughput. In each of the example evolutions located along the maximum number of states boundary in Figure 10.3, the queue distributions have all had the greatest variety of item types. That is, items corresponding to all magazines are present in some numbers in the queue. The absolute fraction of item types is important to the number of states entered and the throughput, even if maximum variety exists, and can significantly affect these attributes, depending on the connectivity of a configuration. For example, for the sparsely connected configuration 94495, with connectivity shown in Figure 10.7, evolutions are identical when the most common item type corresponds to the first, second, or third magazines and have a low number of states (and throughput) because, although the third carriage enters a variety of states, the remaining carriages all halt when the third carriage is transporting items from the first queue to the first magazine (a constant because of de facto priority rules). However, when the most common item type corresponds to the fourth magazine and variety in the queue distribution is still maximal (a (20-20-20-40) queue distribution), the sparsely connected first and second carriages remain operational longer. And although these carriages do not enter any more unique individual states than in other evolutions with maximum variety in the queue distribution, the longer operational times provide sufficient time to allow the third carriage to change it individual pattern, resulting in additional system states.

The association between maximum variety in the queue distribution and the number of unique system states is intuitive. More item types result in potentially more magazines that a carriage can enter, and the loading, transit, and unloading states that follow. But the extent to which additional states are entered is largely dependent on system connectivity, meaning that, although the greatest item variety is a necessary condition, it is insufficient for resulting in the maximum (global) number of states for a system size because suitable connectivity must be present to take advantage of the item variety. However, the greatest item variety typically results in the greatest number of states possible for a particular configuration, although possibly only for specific ratios of item types.

The relationship between the item variety, connectivity, and the number of states (and therefore throughput) is evident in Figure 10.9, which presents a comparison of the distribution of complete evolutions corresponding to the four queue distributions in which all item types are present to the distribution of all complete evolutions. Figure 10.9 reveals that evolutions with the greatest item variety tend to have the greatest number of states at any given throughput. But Figure 10.9 raises some questions, given that the number of states varies considerably at any particular throughput, implying that the correlation between variety in the queue distribution and the number of unique system states is weak and that queue distributions with less variety (absent item types) result in evolutions that enter more states than evolutions corresponding to queue distributions with the greatest variety. While this condition is certainly true when we compare evolutions of different configurations (because connectivity and queue distributions collectively determine an evolution), comparisons of evolutions of any particular demonstrate that item variety does generally correspond to a greater number of states. As an example, consider configuration 127951, which corresponds to the point in Figure 10.9(b) with the fewest number of unique states at near maximum throughput (49 states and $R = 0.658$).

Figure 10.10 shows the distribution of evolutions of configuration 127951 for all possible queue distributions, where highlighted evolutions are those corresponding to the four queue distributions with the greatest item variety. Two of these queue distributions ((20-20-20-40) and (20-20-40-20)) result in the fewest number of unique states at near maximum throughput (the upper left vertex in Figure 10.9(b)), while the remaining two ((20-40-20-20) and (40-20-20-20)) result in the greatest number of states (72 and 64, respectively) and relatively high throughput (0.591 for both) for this configuration. The maximum number of states entered by configuration 127951 is small when compared to all complete 2-3-4 evolutions because the connectivity tends to result in a few global patterns, in which carriages repeat the same sequences, while local patterns dominate in evolutions corresponding to the maximum number of states. However, for configuration 127951, the greatest item variety tends to result in evolutions with the greatest number of states and throughput and this relationship is typical for all configurations, with variations dependent on the absolute ratios of item types taken in the context of the connectivity.

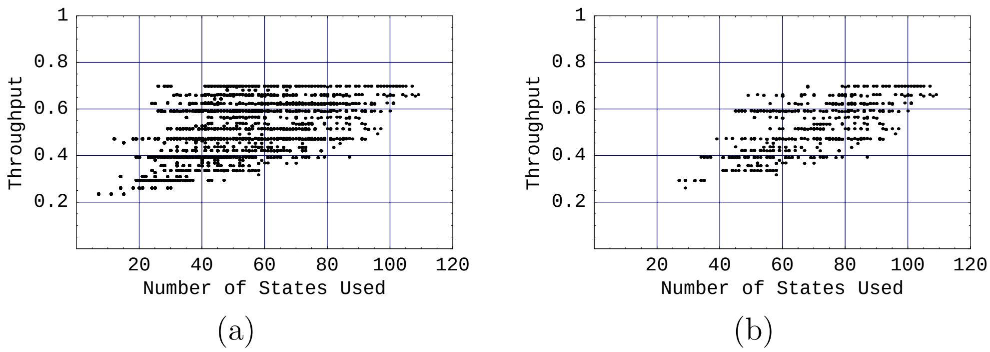

Figure 10.9: A comparison of the distributions of (a) all complete 2-3-4 evolutions and (b) complete evolutions resulting from queue distributions with maximum variety shows that maximum queue variety tends to increase the number of unique states entered in an evolution. However, the number is dependent on the connectivity of the configuration.

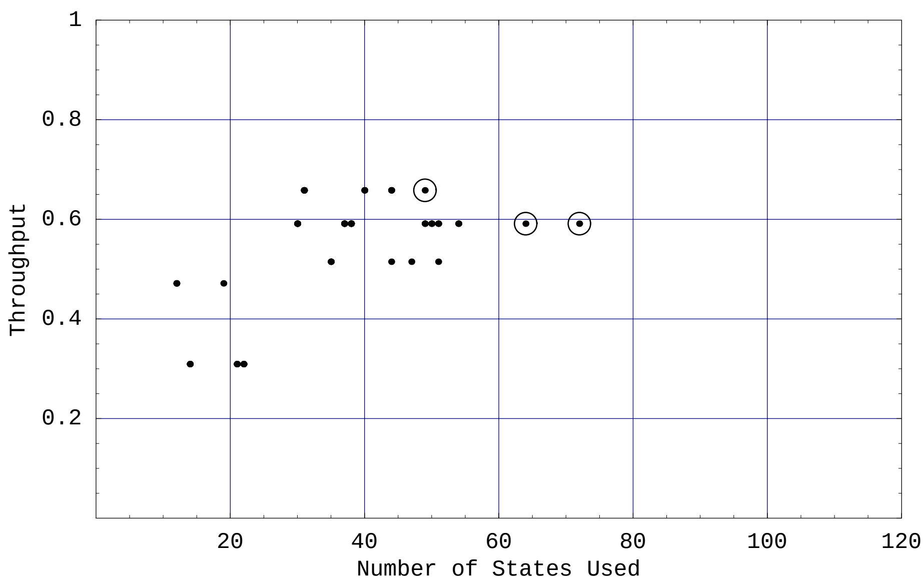

Figure 10.10: The distribution of evolutions for configuration 127951 2-3-4. The highlighted points are those corresponding to evolutions of queue distributions with the greatest item variety. While maximum item variety does not yield the greatest number of states when compared to all complete 2-3-4 evolutions, maximum item variety does result in the maximum or near maximum number of states (and maximum or near maximum throughput) for a particular configuration.

So Figure 10.9 illustrates the effective bounds on the number of unique states and throughput for all configurations given the greatest item variety exists in the queue distribution. These limits are also evident in distributions of sets of evolutions corresponding to progressively lower variety of item types. Figure 10.11 presents the distributions of evolutions with zero to three item types absent and indicate that, as variety decreases and configurations essentially mimic systems with fewer magazines, the minimum and maximum boundaries on the number of unique states in an evolution decreases.

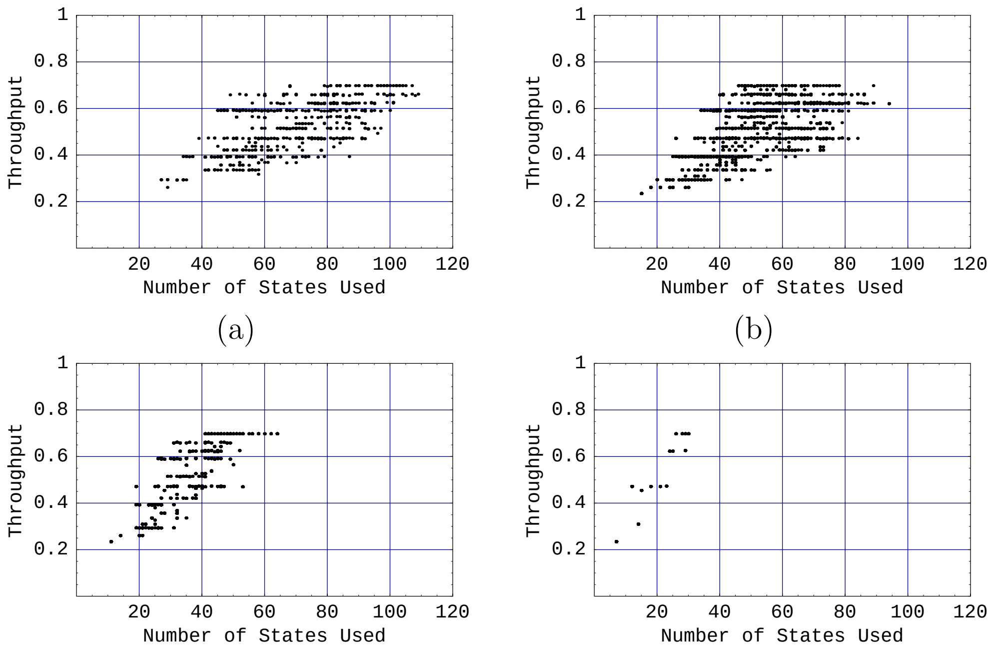

Figure 10.11: The effective bounds on the number of unique states in an evolution for (a) the greatest variety of item types, (b) when a single item type is absent, (c) when two item types are absent, and (d) when a single item type is present in a queue distribution. The absolute location of an evolution is dependent on how the queue distribution 'matches' the connectivity of a configuration, but the figures indicate that as variety of item types increases, the potential for more system states also increases.

Just as the greatest item variety in the queue distributions is a necessary, but insufficient condition for a system to enter the greatest number of unique states in the course of an evolution, Figure 10.11(d) shows that the lowest item variety (those queue distributions consisting of a single item type) is a necessary but insufficient condition for a system entering the fewest number of states. And again, this insufficiency demonstrates that the queue distribution must be considered in the context of the connectivity and shows that the queue distribution and connectivity collectively create a relationship between the number of unique states and throughput.

When only a single shaft is connected to the magazine corresponding to the lone item type and both queues (a necessary condition for a complete evolution), then the system behaves as the simple one carriage, one magazine system with the lowest possible throughput. The only difference between systems under this operating scenario is a function of the number of queues, each of which adds two unique states per pattern, corresponding to travel and loading in each queue. If connectivity is sufficiently complete with respect to the lone item type, then all carriages are involved throughout the evolution and throughput is maximal, with a value approximately equal to the product of the number of shafts and the throughput corresponding to the simple one carriage, one magazine system. The throughput for the simplest system with a single item type for the 2-3-4 systems is 0.235 and the throughput for a sufficiently connected configuration with a single item type is 0.699.

Determining the number of states when all carriages are operational and a single item type is present is more difficult than for the simplest system with a single carriage, since de facto priority logic, resource limitations, and transients are significant factors in an evolution. However, with a single item type and multiple operational carriages, the phase lags that must be present (assuming loading and unloading of one carriage at any queue or magazine at any given time) alone result in more states than in a single carriage evolution.

From the relationships between the number of states and throughput at specific degrees of item variety, and from the conclusions regarding the correlating effects of connectivity on the number of unique states and throughput, we can put together a qualitative relationship between connectivity, item variety, the number of unique states in an evolution, and throughput, which is illustrated in Figure 10.12.

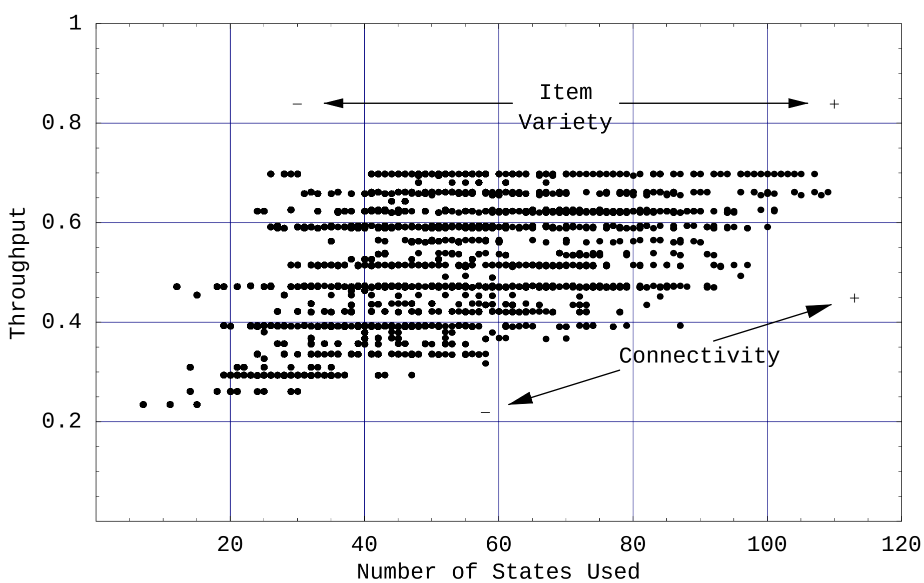

Figure 10.12: Based on the distribution of evolutions with various degrees of variety in the queue distributions in Figure 10.11 and the correlating effects of connectivity on the number of unique states and throughput, we establish qualitative relationships between connectivity, item variety, the number of unique states, and throughput.

Note that the qualitative effect of connectivity is sloped, following our determination that increasing connectivity tends to increase both the number of states and throughput. The slope of the 'connectivity effect' is purely qualitative and an actual slope is an aggregate measurement and will vary across different regions of the number of states/throughput space based on evolution specific dynamics.

The 'item variety effect' is also sloped, indicating that item variety also affects both the number of states and throughput. While it is clear from the distributions in Figure 10.11 that item variety is strongly related to the number of unique states, the distributions do not themselves indicate any relationship between item variety and throughput since they are two-dimensional and the ranges of throughput values are constant for all sets of evolutions with various degrees of item varieties. However, when we look at the mean throughput values for evolutions with different item variety, which are presented in Table 10.1 along with the mean number of states, we see that the mean throughput tends to increase when more item types are present in an evolution. Furthermore, there is a strong correlation (0.988) between the mean number of states and the mean throughput across different degrees of item variety - a relationship that is presented graphically in Figure 10.13.

Table 10.1: The mean number of unique states and throughput as a function of the amount of variety in the queue distribution along with the mean value for all complete evolutions. The mean number of unique states and throughput both increase with greater item variety.

	mean No of States	mean R
Greatest Item Variety	66.2	0.511
Single Item Type Absent	47.5	0.487
Two Item Types Absent	29.9	0.459
Single Item Type	14.0	0.419
All Complete	36.2	0.467

The reason behind the direct relationship between item variety and throughput is related to the probability of carriage utilization. When more variety is present, there is a greater chance for a carriage to be utilized to some extent, and for the evolution to be unique, regardless of the connectivity of its corresponding shaft. That is, if a carriage operates in a niche environment (has limited connectivity to magazines), it will only be utilized when certain item types are present. When more item types are present, the implication is that there is a greater probability that the carriage is utilized to some extent, thereby decreasing the load on other carriages and increasing the throughput. This relation also implies that unique evolutions should tend to have higher throughput - a trend we have seen previously.

As with the qualitative description of the effects of connectivity on the number of states and throughput, the slope of the 'item variety effect' is unclear. However, based on the slopes of curves in the individual distributions of evolutions at various degrees of item variety in Figure 10.11 and the slope of the curve describing the relationship between the mean number of states and throughput corresponding to various degrees of item variety, the variety of item types appears to have a more significant effect on the number of states than on throughput

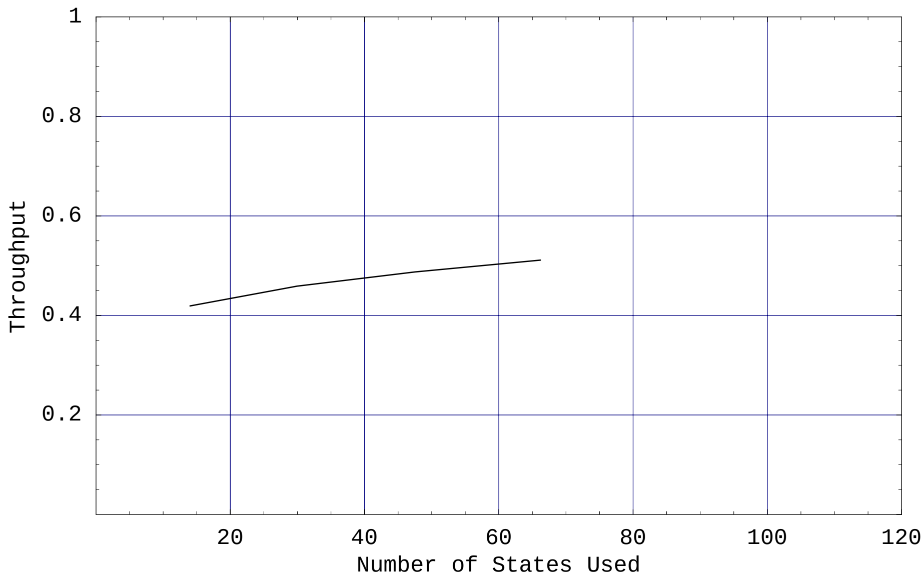

Figure 10.13: The mean number of unique states and throughputs at the four levels of item variety in the queue distributions for 2-3-4 size systems are highly correlated, at a value of 0.988.

and is therefore drawn with a shallower slope than that for the connectivity.

With an understanding of the relationship between connectivity, the number of unique states, and throughput, it is apparent that greater connectivity offers greater adaptability and therefore performance in the face of variable and unknown queue distributions. When design constraints are imposed, specifically practical limitations on physical connectivity, this understanding can result in the identification of configurations with sufficient connectivity to avoid niche behavior by carriages for various queue distributions that leads to halting and decreased performance.

For a fixed configuration - a situation common in many transportation systems, and definitely in an aircraft carrier, where spaces like shafts and magazines can not be moved and paths re-routed - an understanding of the relationship between item variety, the number of unique states, and throughput in the context of connectivity permits the identification or design of suitable queue distributions where the fraction and absolute quantity of item types is matched to the connectivity such that utilization of the shafts/carriages is approximately uniform (implying the evolution is unique as all carriages are utilized) and throughput is maximal.

10.3 Robustness

In the design of configurations, we have shown that the number of unique states is related to the throughput via physical connectivity and can therefore be used as a tool for refining the set of candidate configurations with respect to performance. However, throughput is only one factor in identifying a suitable configuration. We also need to consider the robustness of a configuration and therefore need to identify whether we can use the number of unique states as a tool for identifying candidate configurations with respect to robustness in the same manner as with respect to throughput. That is, is there some correlation between the number of unique states in an evolution and robustness? Certainly, an indirect relationship exists - we know that the number of unique states is related to connectivity and, from the discussion of compressed state complexity, identified relationships between connectivity and robustness. These relationships suggest that some correlation is present between the number of states and robustness which is a function of connectivity.

Figure 10.14 presents the mean number of unique states at various levels of robustness for our example 2-3-4 size systems, normalized by the maximum mean number of unique states. As for figures illustrating the relationship between mean complexity measures and robustness in the discussion of compressed state complexity, levels of robustness with no corresponding evolutions, and therefore a mean number of states equal to '0', are excluded to create mean complexity 'curves' as a graphical device which do not imply that non-integer levels of robustness exist or that mean values should exist at levels of robustness that are not populated with any evolutions.

The shape of the normalized mean number of states curve is quite similar to the state complexity and compressed state complexity curves and, to a lesser extent, the logical complexity and throughput corves presented in the discussion of compressed state complexity. The similarity in shapes indicates not only a relationship between the number of unique states and other complexity measures (which is no surprise with respect to the state and compressed state complexities which by definition are directly related to the number of states), but also provide further evidence of a correlation between the number of states and throughput. The similarity in the shapes of the curves also implies that any conclusions drawn from the curves of other complexity measures also apply to the number of unique states - in particular, we are interested in the reasons behind trends in the curves resulting from connectivity. That is, how connectivity is related robustness and therefore how robustness and the number of states are related through a common association with connectivity. We can draw conclusions regarding a correlation between the number of states and robustness directly from the normalized curve in Figure 10.14 in the context of the populations of configurations at each level of robustness, as well as from the curve of numbers of states of the set of unique, complete configurations.

Figure 10.15 presents the same curve relating the mean number of states over the range of robustness values, but superposed with the normalized number of configurations at each

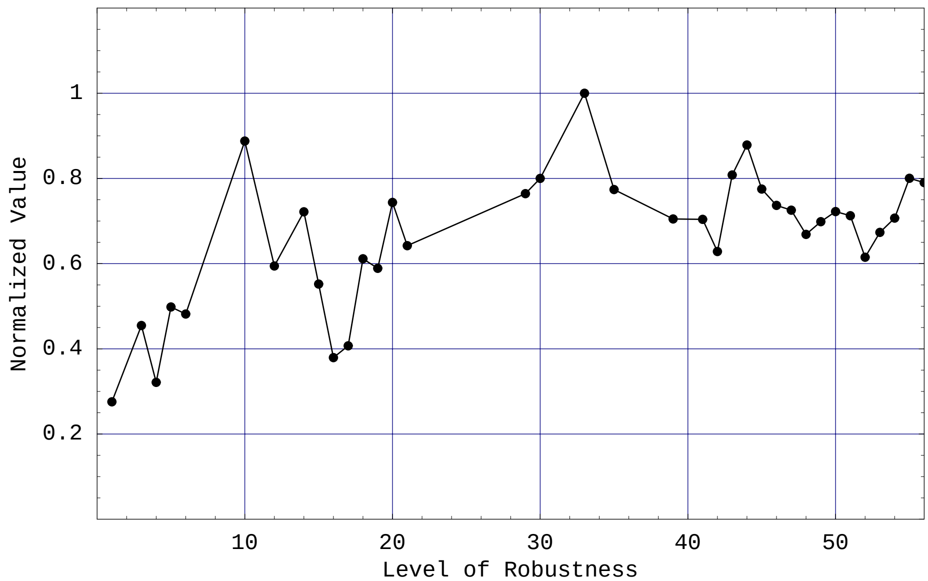

Figure 10.14: The normalized mean number of unique states 'curve' across possible levels of robustness. The shape of the curve is nearly identical to the curves for state complexity and compressed state complexity, implying any conclusions drawn from those curves also apply to the number of unique states.

level of robustness. Evolutions must be complete for configurations with present connectivity for a particular number of magazines with respect to queues when the queue distribution is composed solely of items corresponding to those magazines. For queue distributions containing items corresponding to magazines with absent connectivity to queues, the evolutions can only be incomplete. Therefore, for the set of complete evolutions, most configurations exist at the most robust levels of robustness for various numbers of magazines (the first level for a single magazine with present connectivity to all queues, the sixth level for two magazines with present connectivity, the 21st level for three magazines with present connectivity, and the 56th level when all magazines have present connectivity). A configuration will only have an intermediate level of robustness if halting occurs for some queue distributions. Because the number of configurations at levels of robustness corresponding to robust mimics are significantly greater than at intermediate levels, the mean complexity values of these dominant populations are comparable since the amount of averaging is on the same order of magnitude and offer a more meaningful comparison and correlation than when we consider the mean number of states at all levels of robustness. So, while the correlation between the mean number of unique states across all levels of robustness is 0.599, the correlation between the mean number of unique states and levels of robustness corresponding to the various dominant levels of robust mimicry, which are highlighted in Figure 10.15, is significantly greater at 0.907. This strong correlation for the more meaningful, refined set of configurations suggests that configurations with evolutions that enter, on average, a greater number of unique states, also tend to demonstrate adaptability with respect to variations in the queue distributions.

Furthermore, since we know that configurations with some amount of absent connectivity between queues and magazines can be complete only for a subset of queue distributions and are therefore less robust and that a complete evolution of a configuration with absent connectivity involves fewer magazines and therefore has a smaller potential state space, we see that the relationship between robustness and the number of states is a result of connectivity.

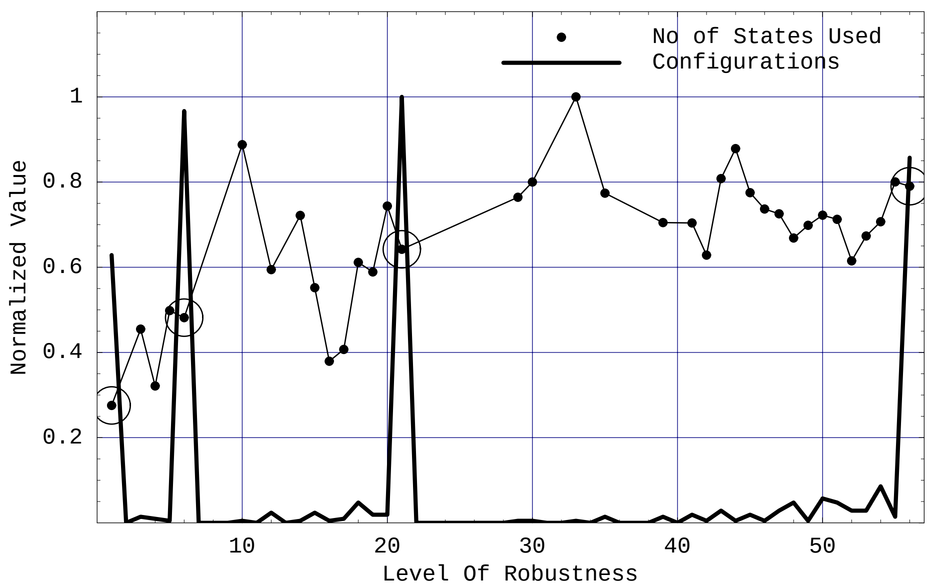

Figure 10.15: The normalized mean number of unique states at all levels of robustness, superposed with the number of configurations comprising each level for the set of 2-3-4 evolutions. The 1st, 6th, 21st, and 56th levels - those corresponding to the most robust mimics - dominate the population and comparison of the highlighted mean numbers of unique states at these levels is more meaningful, since the numbers of configurations at these levels are on the same order of magnitude. When we account for populations and comparable averaging, a correlation between the number of unique states and robustness is apparent.

Since the set of all complete evolutions contains evolutions that mimic evolutions with fewer than the nominal number specified by the configuration size (when not all carriages are utilized), the refined set of unique complete evolutions arguably presents a more accurate representation of the capabilities of the nominal system size. Figure 10.16 presents the relationship between the normalized mean number of unique states over the range of possible robustness values, superposed with the normalized configuration populations. The populations of configurations are no longer concentrated at the levels corresponding to the most robust mimics as in Figure 10.15 for all complete evolutions, suggesting a more meaningful correlation between the mean number of states and robustness exists when all levels of robustness are considered.

The relationship between the mean number of unique states and robustness for the set of unique evolutions is apparent in Figure 10.16, with a correlation value of 0.836. As with the

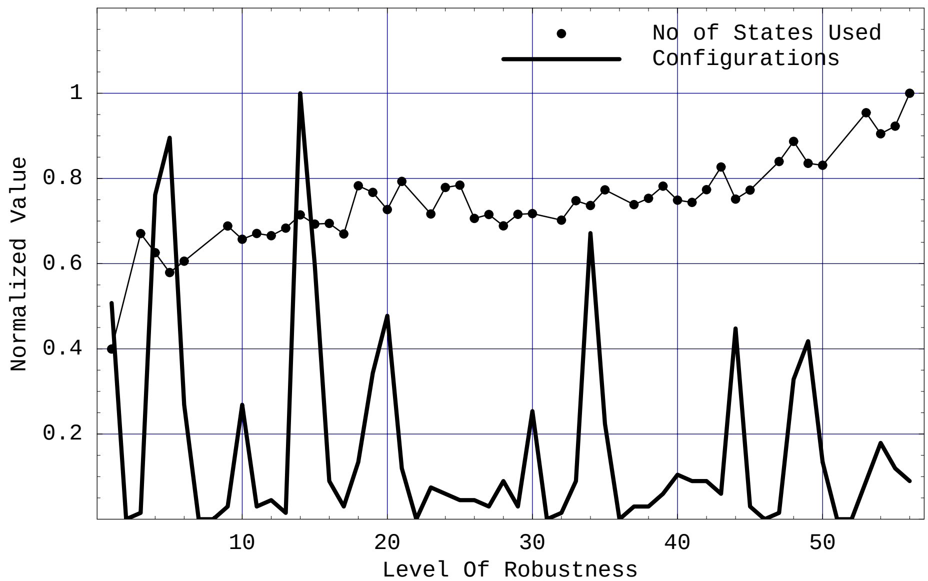

Figure 10.16: The normalized mean number of unique states at all levels of robustness, superposed with the number of configurations comprising each level for the set of unique 2-3-4 evolutions. The 1st, 6th, 21st, and 56th levels - those corresponding to the most robust mimics - dominate the population and comparison of the highlighted mean numbers of unique states at these levels is more meaningful, since the numbers of configurations at these levels are on the same order of magnitude. When we account for populations and comparable averaging, a correlation between the number of unique states and robustness is apparent.

set of all complete evolutions, we again see evidence that connectivity is the link between the mean number of states and robustness. Just as for the set of all complete evolutions, the amount of absent connectivity between queues and magazines determines the number of queue distributions that result in complete evolutions, and therefore the level of robustness, for the set of complete, unique evolutions. Since only queue distributions lacking item types corresponding to magazines with absent connectivity can result in complete evolutions, lower connectivity results in lower robustness because there are fewer deliverable queue distributions and fewer item types result in fewer system states because carriages have fewer possible destinations, resulting in a correlation between the number of system states and robustness via connectivity.

However, when we consider only unique evolutions, the mean number of states for any configuration should be greater than (or equal to) the mean for that configuration when all complete evolutions are taken into account because the involvement of all possible carriages tends to increase the potential state space. Furthermore, the number of complete unique evolutions corresponding to a given configuration should be less than (or equal to) the number of all complete evolutions for that configuration, depending on the number of evolutions of that configuration that do not involve all carriages, which is a function of connectivity. Therefore, for unique complete evolutions, the level of robustness of a given configuration is less than (or equal to) the level of robustness when all complete evolutions are considered, meaning that lower levels of robustness should have greater populations, but greater mean numbers of states. The direct relationship between the mean number of states and robustness should still remain however, since the number of states for an evolution involving a certain number of magazines should always be greater than an evolution involving fewer magazines.

In the analysis of the number of unique states as a complexity measure, we have shown that, although relationships between the number of states and throughput are present, the number of unique states as a complexity measure does not violate the definition of algorithmic complexity, which does not exhibit any relationship to performance. The inherent relationships between the number of unique states and throughput is strongly related to connectivity, which when placed in the context of the amount of item variety on queue distributions, leads to a qualitative mapping of the number of states/throughput space. Connectivity was also shown to play an important role in the relationship between the number of unique states and robustness, for both all complete and unique complete evolution sets. The importance of physical connectivity to various complexity measures, robustness, and throughput leads us to examine the use of connectivity as a predictive measure of complexity, performance, and adaptability in the following section.

Chapter 11

Static Complexity Measures

In previous sections, we have seen evidence that physical connectivity is the attribute that plays a significant role in the correlations seen between several complexity measures and throughput. Physical connectivity is also an important factor in defining the adaptability of a configuration and therefore in the relationships between various dynamic complexity measures and robustness. This evidence would suggest that physical connectivity alone, which is expressible as a static complexity measure, could be related to throughput and robustness to the point where searches for optimal configurations could occur without the expense of explicit simulations.

While evidence of the value of physical connectivity as a static measure exists, we begin by examining the utility of other static measures: the total number of possible system states, average logical connectivity, and the fraction of potential physical connectivity.

11.1 Total Number of Possible States

The total number of possible system states for a configuration is primarily a function of system size, physical connectivity of system elements, and other physical attributes, such as resource availability and the number of carriages that can simultaneously occupy a given space. These additional physical attributes are not scalable, and are assumed to remain constant for all system sizes.

The total number of possible states does not account for specific queue distributions so, even though a carriage may never enter a magazine because no corresponding item types are present, states involving that carriage and that magazine are still included, meaning the total number of possible states for a specific queue distribution is often an overestimate. However, under the assumption that queues are uncontrollable, we should expect that all queue distributions are possible and all possible states will occur in the life of the system.

Since the total number of possible states is intended as a potential measure, not a measure of ranges of possible values of a configuration, each configuration corresponds to a single number of possible states, meaning the size of the set of values of possible states is always a fraction of the size of the set of values for any measure that accounts for specific queue distributions ($\frac{1}{56}^{th}$ for systems with four magazines, $\frac{1}{21}^{st}$ for systems with three magazines, etc...).

Figure 11.1 presents the two-dimensional distributions of the total number of possible states with respect to throughput for several system sizes and helps reveal the relative effects of the number of queues, shafts, and magazines on the potential state space. More importantly, the figures reveal that no apparent relationship between the total number of possible states and throughput exists for any system size, which is supported by weak correlation values of approximately 0.35 for most system sizes.

Figure 11.2 presents the robustness "curve" of the normalized mean number of possible states for all complete evolutions of configurations with 2 queues, 3 shafts, and 4 magazines, superposed with normalized populations of configurations at each level of robustness and highlights of mean values at levels of robustness with comparable populations that result in more meaningful aggregate values. Even though the number of possible states is a purely static measure, the shape of the robustness curve remains similar to the curves for dynamic measures related to the number of unique states entered in the course of an evolution, and to a lesser extent, the throughput and logical complexity curves. The highlighted values at the 1^{st}, 6^{th}, 21^{st}, and 56^{th} levels of robustness exhibit a more linear relationship with respect to robustness than these other curves, which raises the corresponding correlation value to 0.994. When we consider unique complete evolutions, the relationship of the mean number of possible states to robustness is again similar to those of dynamic measures, which is evident in Figure 11.3. The mean number of potential states again increases with robustness and, although the potential states/robustness curve is more jagged than the equivalent dynamic complexity measures, exhibits a correlation value of 0.865.

When looking at the similarities between robustness curves for the total number of possible states and other dynamic measures, the question arises of why no apparent relationship exists between the number of possible states and throughput in Figure 11.1, especially when we consider Figure 11.4, which indicates a strong relationship between the number of potential states and the average physical connectivity, since physical connectivity has been shown as a significant factor in introducing correlations between dynamic measures involving numbers of states and throughput.

The answer lies partially in the fact that the total number of possible states is only partly a function of connectivity. Even for a fixed system size, the total number of possible states is still a function of other system attributes, like resource allocation and assumed space limitations. More importantly, the apparent contradiction is a result of the fact that the potential number of states is a static measure that does not account for variations in queue distributions and is constant for a configuration, even though a configuration may result in

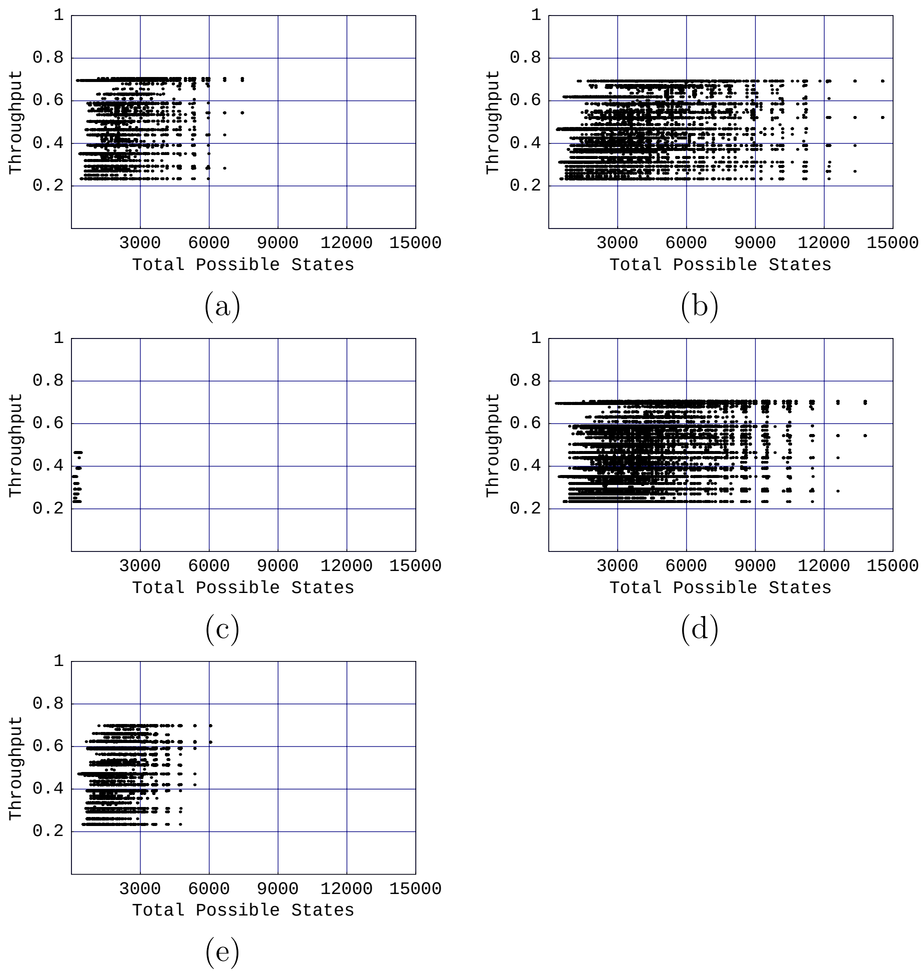

Figure 11.1: The distribution of total possible states with respect to throughput for complete evolutions for (a) 3-3-3, (b) 4-3-3, (c) 3-2-3, (d) 3-3-4, and (e) 2-3-4 size systems. The distributions use the range of throughput values possible for a given configuration, which has a constant total number of possible states, and therefore reveals no apparent relationship between possible states and throughput.

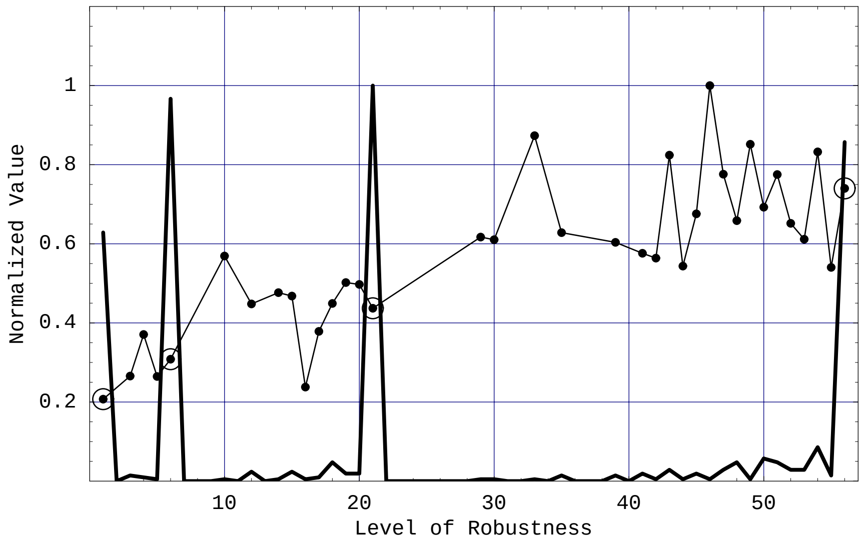

Figure 11.2: The robustness curve for the total number of possible states for 2-3-4 size systems is similar to those for dynamic measures related to the number of unique states entered. When the mean numbers of potential states corresponding to the most populated levels of robustness are considered, the correlation value between the total number of states and robustness is 0.994.

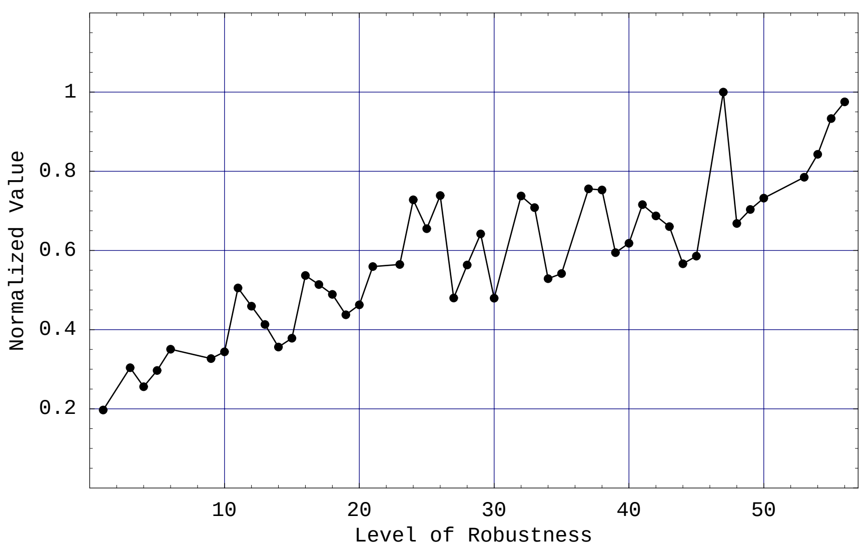

Figure 11.3: The robustness curve for 2-3-4 size systems with respect to the number of possible states for complete unique evolutions indicates a direct relationship between the potential state space and robustness. The correlation value between the number of potential states and robustness for complete unique evolutions is 0.865.

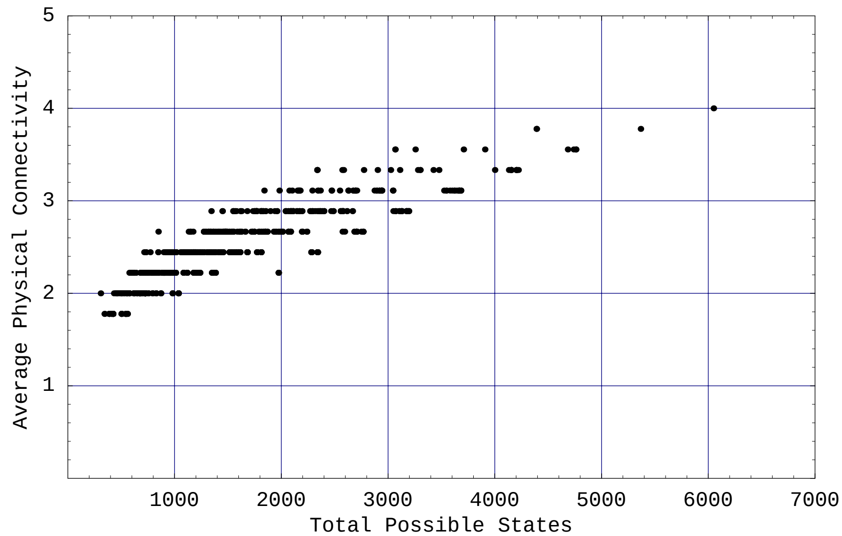

Figure 11.4: The total possible states and average physical connectivity of complete evolutions of 2-3-4 system size have a correlation value of 0.940. This relationship suggests that the total number of possible states is also related to dynamic measures involving the actual number of states entered in an evolution.

a range of throughputs in response to different queue distributions.

The effect of ignoring variations in queue distributions is partly reduced by using mean values - either mean throughput values or mean values of possible states at any level of robustness. The use of aggregate values helps explain why the robustness curves for the potential numbers of states are similar to those for dynamic measures, which also utilize average values, although the mean values for dynamic measures account for variations internal to each configuration. When the average throughput values for each configuration are considered rather than explicit throughput values, the correlation between the number of potential states and throughput increases to values approaching those for dynamic measures. Figure 11.5 presents the distribution of mean throughput values with respect to the number of potential states, where each point now represents the attributes for a unique configuration. The distribution has a correlation value of 0.458 and a definitive maximum number of potential states boundary. Conceptually, the use of aggregate values for potential complexity measures is slightly more intuitive. With static measures in particular, we are interested in the potential behavior of a configuration rather than values specific to any particular evolution.

11.2 Average Logical Connectivity

Unlike other complexity measures, average logical connectivity is not based on any attributes specific to an evolution or configuration. Average logical connectivity is intended to quantify

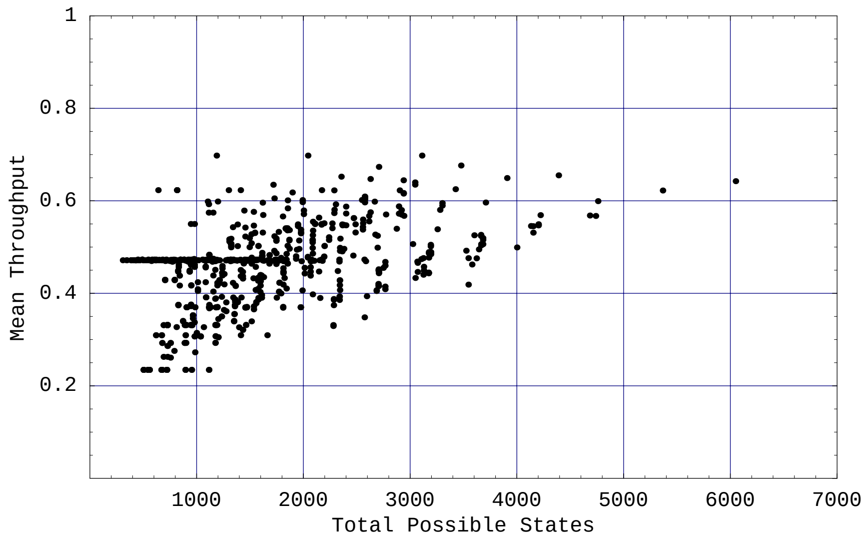

Figure 11.5: The distribution of the 2-3-4 configurations with respect to their mean throughput for all queue distributions that result in complete evolutions and the total possible system states. Use of the mean throughput for all queue distributions that result in complete evolutions provides an indication of the range of performance with respect to a given number of possible states.

the complexity of the system rules, which is a constant for all configurations of a given system size. Just as a complete description of the total number of possible states yielded a range of throughput values for a single configuration, a complete description of the average logical connectivity provides a range of throughput values for a single system size. The average logical connectivity is therefore most suited for relative comparisons between systems with different numbers of queues, shafts, and magazines.

While it is fairly straightforward to determine the range of throughput values possible for any given system size, the distribution of throughput values is typically not obvious. As with the total number of possible states, a simple, first-order method of describing the distribution is with the mean throughput, now associated with a particular system size. Figure 11.6 presents the distribution of systems with respect to their average logical connectivities and mean throughput values for *(a)* all system sizes and *(b)* "interesting" system sizes, where interesting system sizes are those in which non-trivial behavior occurs. For instance, systems with a single shaft consist of a single configuration that has constant performance regardless of the queue distribution and are considered trivial. The correlation values between the average logical connectivity and mean throughput are 0.518 and 0.365 for all system sizes and interesting system sizes, respectively. Although the mean throughput for all "uninteresting" system sizes are the lowest possible, since many represent the simple single carriage systems, the difference in correlation values suggests that the specific members of the set of system sizes are very important in establishing any relationship between logical connectivity and performance as the omission or inclusion of a particular system size can have a significant impact on any relationship.

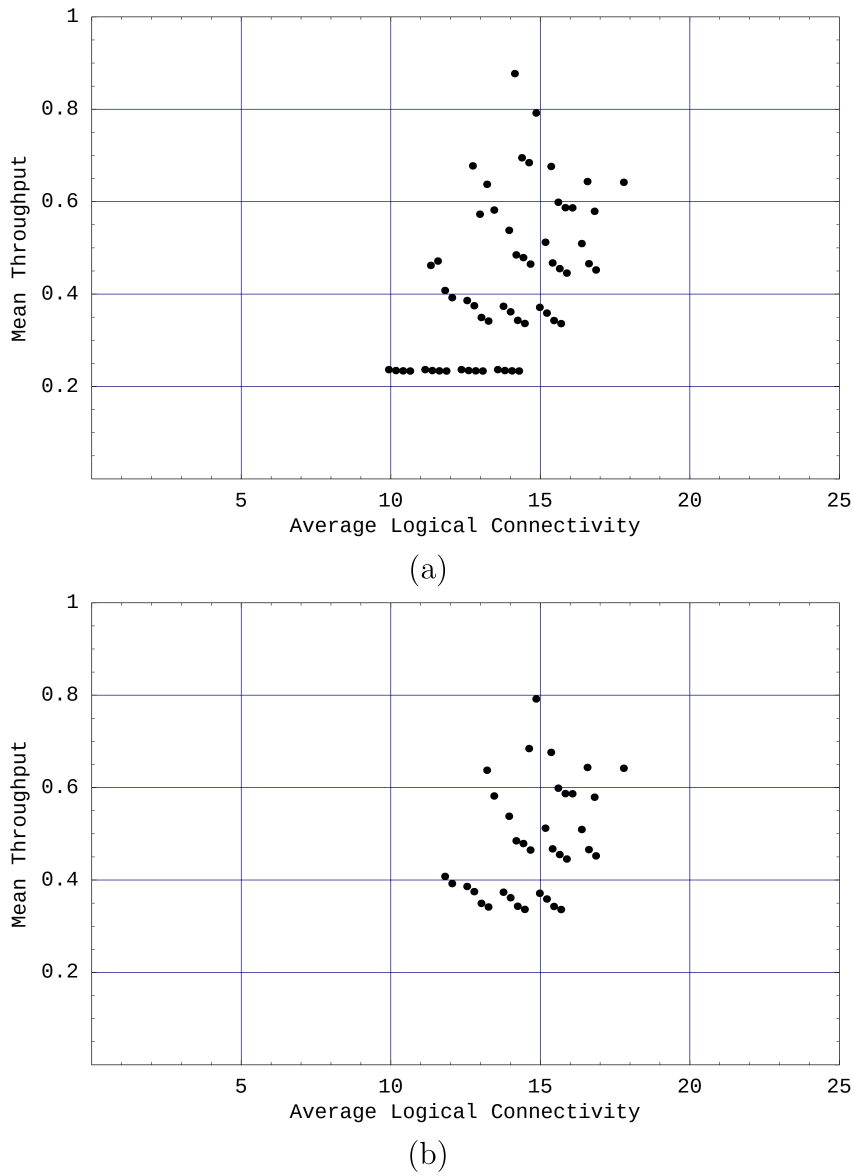

Figure 11.6: The distribution of average logical connectivities and mean throughputs for (a) all system sizes considered and (b) "interesting" systems sizes that contain non-trivial evolutions.

Apparent relationships between the average logical connectivity and throughput, and any resulting conclusions, can be flawed for additional reasons. Comparisons of attributes across different system sizes are not particularly valid, especially when only mean values are considered, because the populations of various system sizes are typically not comparable. Additional flaws with average logical connectivity are more fundamental in nature. Larger system sizes result in more information involved in the application of evolution logic, which is reflected in the average connectivity. But does more information involved in a logic rule necessarily result in greater complexity? According to Kauffman, greater connectivity results in systems that tend towards chaos, but behavior is a function of the *particular* logical connectivity of a system, regardless of its size. While average logical connectivity indicates the amount of connectivity, which is not near complete and scales with system size depending on the absolute and relative numbers of each element type, it does not tell us anything about the particular rules, limiting the effectiveness of average logical connectivity as a complexity measure.

Furthermore, since average logical connectivity is a constant for a given system size, we are unable to distinguish between simple and complex behaviors that occur in evolutions of that system size. To use logical connectivity as a complexity measure, it would have to assume the form of an evolution or configuration specific measure, such as rules employed in the course of an evolution, which no longer results in a static measure. Such a measure is close, but not identical, to the number of unique states entered in an evolution, since the number of unique states entered does not consider state trajectories. State trajectories are not necessarily unique since not all information is included in state definitions.

Although potential exists for the use of logical connectivity as a measure of complexity, a static measure involving average values for a system size has several shortcomings and has only limited use in the relative comparison of systems of various sizes. In this role, average logical connectivity presents nothing new - primarily that larger systems are capable of supporting greater complexity. However, it does this by indicating that more information is required to evolve larger systems.

11.3 Average Physical Connectivity

For all configurations of a given system size, the set of queue distributions considered for each configuration is constant. As a result, it is reasonable to assume that any relationship between complexity and performance evident for a set of evolutions of a given system size is largely the result of physical connectivity, since connectivity is the only variable attribute that distinguishes configurations. And we have already seen evidence of the role connectivity plays in relationships between complexity and performance and between complexity and adaptability in the discussions of dynamic complexity measures, particularly those related to the number of unique states. Physical connectivity therefore appears to offer great potential as a useful static complexity measure that significantly reduces the cost associated with explicit simulation in the design of optimal configurations.

To express the physical connectivity of a configuration as a single value, the average connectivity for each element in the configuration is used. While the single average value is a useful format as a complexity measure, the information describing the explicit definition of a configuration, which we have seen can have a significant impact on performance and adaptability, is lost. Configurations that have identical average connectivities may therefore correspond to a range of throughput and robustness values.

As with the total number of possible states, the average physical connectivity is a configuration specific measure, not an evolution specific measure. As a result, the distributions of evolutions with respect to average physical connectivity and throughput, like the example distribution in Figure 11.7 for 2-3-4 size systems show no apparent correlations because each configuration corresponds to a range of throughput values and a single physical connectivity value.

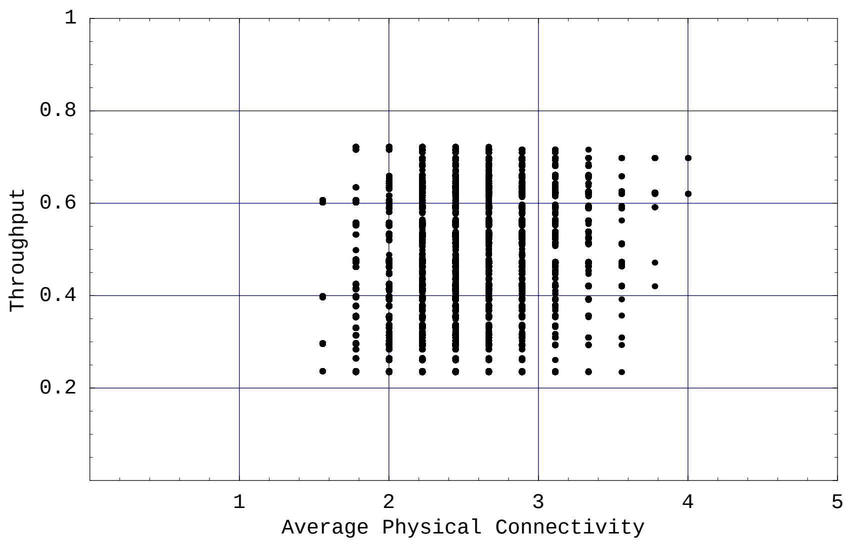

Figure 11.7: The two-dimensional distribution of complete 2-3-4 evolutions with respect to average physical connectivity and throughput reveals no apparent relationship because average physical connectivity is a configuration-specific attribute.

But we also saw with the total number of possible states that mean throughput values are useful for identifying relationships for configuration-specific complexity measures because the mean provides a simple, first-order measure of potential performance. Figure 11.8 presents a comparison of mean throughput values for all configurations that result in at least one complete evolution. A direct relationship between the average values of connectivity and throughput is evident in Figure 11.8 although the correlation value is relatively low at 0.572.

The results of halting evolutions are not included in the characterization of average physical connectivity simply because the throughput values associated halting evolutions results in a misleading characterization of the mean throughput for the corresponding configuration. Following similar logic, only configurations that result in complete evolutions for the same set of queue distributions are comparable because average connectivity is defined to consider the connectivity of *all* elements, regardless of their relevance with respect to the set of queue distributions that correspond to a complete evolution.

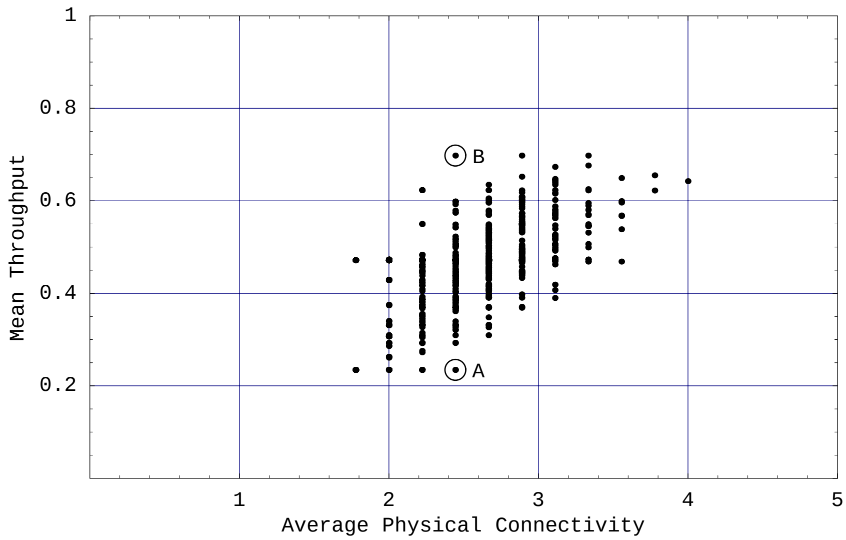

Figure 11.8: The distribution of 2-3-4 configurations with respect to their average physical connectivities and the mean throughputs of all evolutions that are complete for a given configuration. The highlighted configurations are used to illustrate the important distinction between average physical connectivity of an entire configuration and the average physical connectivity of relevant elements.

For the same set of queue distributions, comparisons of average physical connectivities are valid when all configurations result in complete evolutions for all queue distributions in the set because each mean represents a complete characterization of the performance for that set. However, for configurations that result in complete evolutions for a fraction of queue distributions because of absent connectivity between queues and magazines, we only get an idea of the performance for that fraction, which is dependent only on the connectivity of relevant magazines (those magazines with present connectivity to all queues and define a configuration that is complete for all queue distributions comprised solely of items corresponding to those relevant magazines). Yet, in calculating average physical connectivity, we consider the connectivity of all magazines (and shafts connected to those magazines), even if they are irrelevant, which results in a misleading representation of relevant connectivity.

To illustrate the effects of considering the connectivity of only relevant magazines, we consider two configurations with average physical connectivities and mean throughputs highlighted in Figure 11.8. Configuration 217582, which has the connectivity presented in Figure 11.9 and corresponds to point A in Figure 11.8, is only capable of complete evolutions for a single queue distribution (0-0-0-100) because only the fourth magazine has present connectivity to all queues. Throughput is low because a single carriage transports all items from both queues to the fourth magazine. Yet the connectivity of the remaining, irrelevant magazines is high (actually, it is the maximum possible for irrelevant magazines for this system size). So the "mean" throughput is low because of the connectivity of the single relevant magazine, but the average connectivity is relatively high because irrelevant magazines are fairly well connected.

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

Figure 11.9: The incidence matrices for configuration 217582 2-3-4 show that only the fourth magazine is relevant. The configuration is therefore only complete for a single queue distribution (0-0-0-100). For this queue distribution, throughput is low because a single carriage delivers all items, yet the connectivity of irrelevant magazines is relatively high.

Conversely, consider configuration 250143, which is represented by point B in Figure 11.8 and has the connectivity illustrated in Figure 11.10. Once again, there is a single relevant magazine, but there are multiple paths between this magazine and each queue, meaning multiple carriages are involved in the evolution in which only item types corresponding to the relevant magazine are present in the queues. The utilization of multiple carriages results in a high throughput. However, despite greater connectivity of the relevant magazine, configuration 250143 shares the same average physical connectivity with configuration 217582 because it has sparse connectivity (the sparsest possible) for the irrelevant magazines.

$$(SQ) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 0 & 1 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 0 & 0 & 0 & 2 \\ 1 & 1 & 1 & 3 \end{pmatrix}$$

Figure 11.10: The incidence matrices for configuration 250143 2-3-4. Like configuration 217582, only the fourth magazine has present connectivity to all queues and the configuration results in a complete evolution for only a single queue distribution. However, the connectivity of the single relevant magazine is high, so throughput for the complete evolution is high. The connectivity of the irrelevant magazines is low, so the average connectivity of the configuration is equal to that for configuration 217582.

So configurations 217582 and 250143 are opposites with respect to their connectivities in terms of relevant and irrelevant magazines, which affects their mean throughput values. Configurations such as these misrepresent the average physical connectivity/mean throughput distribution, not through their throughput values, but because of the connectivity of irrelevant magazines since they have no bearing on the throughput for the set of queue distributions for which the configurations result in complete evolutions.

In order to account for the effects of the connectivity of irrelevant magazines on the average physical connectivity and its distribution with respect to mean throughput, we can take several approaches. We could consider only the connectivity of queues, shafts, and *relevant* magazines in calculating an average physical connectivity. However, this approach would prohibit the comparison of configurations with different numbers of relevant magazines, without some form of normalization. Similarly, we could make physical connectivity an evolution specific measure by considering only those magazines and shafts (all queues must be considered) that are involved in a given evolution, which is a function of both connectivity and the queue distribution. While this approach may provide a better characterization of the relationship between connectivity and throughput, we lose the ability to identify potentially optimal configurations, since a configuration is defined by the connectivity of all components in the context of all queue distributions. Perhaps the most useful approach is to consider only configurations that result in complete evolutions for the same set of queue distributions, in particular, those configurations that result in complete evolutions for all queue distributions. With this approach, the number of relevant magazines must be constant so comparisons of average connectivities are valid. Additionally, in terms of identification of optimal configurations, where adaptability is a critical attribute, we are most interested in these most robust configurations anyway, so omission of configurations that are not the most robust is not detrimental.

Figure 11.11 presents the distribution of evolutions with respect to average physical connectivity and mean throughput for the most robust configurations of 2-3-4 system size. The relationship between average physical connectivity and mean throughput is more defined than that for the set of configurations with at least one complete evolution, and the correlation is significantly higher at 0.720, which suggests that average physical connectivity can be a useful tool as a potential measure of performance. Examining the highlighted configurations in Figure 11.11 in detail enables us to further understand the relationship between connectivity and throughput.

Configuration 110878, which corresponds to point A in Figure 11.11, represents the configuration with the lowest average physical connectivity that is still capable of complete evolutions for all queue distributions. It is interesting to note however, that despite its sparse connectivity, 110878 does not have the lowest mean throughput. The connectivity of 110878, presented in Figure 11.12, results in the lowest possible throughput for any queue distribution that does not include items corresponding to the fourth magazine because the third carriage is the only carriage utilized. For queue distributions that include items bound for the fourth magazine, the throughput improves, depending on how much the first and second carriages are utilized.

Lower mean throughput is possible when shafts have identical sparse connectivity, which induces earlier carriage haltings for queue distributions in which items corresponding to the shared magazine are present. We have already identified such a configuration in the discussion of the number of unique states and it is therefore no surprise that configuration 94495 corresponds to point B in Figure 11.11. The connectivity of configuration 94495 is presented again in Figure 11.13 and shows that, despite the additional cost of additional connectivity, throughput is lower than a more sparsely connected configuration because the identical connectivity of the first and second shafts lowers their potential utilization. Configuration 94495 illustrates the point that the actual connections are important when considering connectivity and the limitations of using average values.

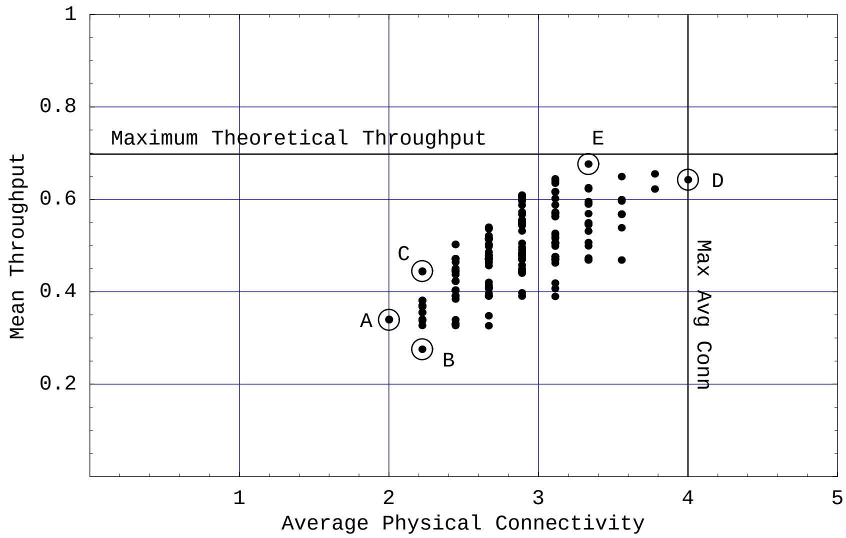

Figure 11.11: The distribution of the most robust 2-3-4 configurations with respect to their average physical connectivities and mean throughputs. Comparison of configurations that result in complete evolutions for the same set of queue distributions presents a valid characterization of the relationship between physical connectivity and throughput. Detailed analysis of the highlighted configurations provides a better understanding of how connectivity and throughput are related.

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

Figure 11.12: The incidence matrices for configuration 110878 2-3-4. Configuration 110878 has the lowest possible average physical connectivity that still results in complete evolutions for all possible queue distributions. However, it does not correspond to the lowest mean throughput.

$$(SQ) = \begin{pmatrix} 0 & 1 \\ 0 & 1 \\ 1 & 1 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 3 \end{pmatrix}$$

Figure 11.13: The average physical connectivity for configuration 94495 is greater than that for 110878, but the corresponding mean throughput is lower because the first and second shafts have identical connectivity with respect to both queues and magazines. Identical connectivity effectively halves the utilization of both carriages and the remaining carriage is always the limiting factor.

Configuration 222156, which corresponds to point C in Figure 11.11, further illustrates the importance of considering the actual connections of a configuration. Configuration 222156 has an average physical connectivity identical to that for configuration 94495. But, like configuration 110878, 222156 has a single path between each queue and magazine, which is evident in the connectivity of configuration 222156 in Figure 11.14. Also like configuration 110878, configuration 222156 has two shafts with identical connectivity with respect to magazines, but distinct connectivity with respect to queues for these same shafts. It is in this regard that 222156 differs the most from configuration 94495 - the same shafts that have identical connectivity with respect to magazines also have the same connectivity with respect to queues in 94495, which is a primary reason for its poor performance. When considered collectively, the behaviors of configurations 94495, 110878, and 222156 suggest that, in the face of constraints on the number of connections possible, the absence of redundant connectivity results in greater overall performance when considering a range of potential queue distributions.

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

Figure 11.14: Like configuration 110878, configuration 222156 has two shafts with identical connectivity with respect to magazines, but dissimilar connectivity with respect to queues for these same shafts. The lack of any redundancy permits greater mean throughputs when connectivity is generally sparse.

Figure 11.11 indicates that, in general, greater physical connectivity results in greater throughput. This trend is intuitive, as more potential paths creates more uniform carriage utilization and delays carriage haltings, which increases throughput. Of course, not apparent in Figure 11.11 is the trade-off in the cost of greater physical connectivity. Additionally, this trend must be placed in the context of the conclusions reached from the comparison of configurations 94495 and 222156 that actual connectivity, particularly the amount of redundancy, is a significant factor with respect to performance.

While redundancy negatively affects throughput when identical shafts have low average connectivity, redundancy tends to improve throughput at sufficiently high connectivity, particularly when the redundant shafts share complete connectivity to all queues. Redundant shafts tend to have the same utilization. Therefore, since uniform utilization for all carriages equates to the greatest throughput, as the connectivity of redundant shafts increases, the number of queue distributions that result in more uniform utilization increases as does mean throughput, although throughput is still a function of the connectivity of distinct shafts. So, for the configuration with complete connectivity, carriage utilization is theoretically the most uniform for all queue distributions and throughput is theoretically maximal because of complete redundancy.

However, Figure 11.11 indicates that the completely connected configuration, represented by point D, does not have maximal throughput. The discrepancy in the theoretical and actual values is not a function of the connectivity, but the applied evolution logic. For the logic used, a carriage selects its destination queue after unloading based on the state of the queue inventories at the time the carriage is in the magazine, not when the carriage is expected to reach the queue. With this logic, it is possible for a carriage to arrive at a queue to find that the item on which it based its decision has been removed by another carriage. If there are no other transportable items for the carriage, the carriage remains in the queue for the remainder of the evolution, effectively halting and resulting in lower than maximal throughput. A detailed illustration of the progression of this logic is presented in Appendix B. This scenario is more common in configurations with greater connectivity and greater redundancy, where carriages share common destinations.

This halting "problem" can be rectified by altering the evolution logic to have a carriage reserve items in a queue, which are then set aside in a buffer to await the arrival of that specific carriage. Using this evolution logic, the theoretical mean throughput for the completely connected configuration is equal to the maximum theoretical throughput for any configuration of the same system size, which is illustrated as an upper boundary in Figure 11.11 - in this case, for the 2-3-4 system size.

Figure 11.11 illustrates an interesting phenomena that can occur at higher average connectivities. The intersection of the sloped line through the locus of configurations that define the maximum mean throughput boundary and the vertical line defining the maximum average physical connectivity do not intersect at the maximum theoretical throughput boundary. Instead, a "knee" occurs when average physical connectivity is sufficiently high, where mean throughput is near maximal. This knee indicates that maximal, or near maximal throughput is possible for all queue distributions for a configuration without complete connectivity. Adding connections beyond this point does not improve throughput significantly, but adds additional practical costs associated with establishing physical connections.

The configuration corresponding to this knee, point E in Figure 11.11, is configuration 251839, with connectivity illustrated in Figure 11.15. The configuration never experiences over-utilization of any carriage due to shaft redundancy. Disproportionate carriage utilization only occurs when there are sufficient item types corresponding to the first or second magazines, because the second carriage must deliver all items from the first queue to the first magazine or the first carriage must deliver all items from the first queue to the second magazine. If control of the queues, sufficient to prevent a large fraction of items bound for the first or second magazines in the first queue, is possible, then configuration 251839 is immune to carriage haltings and will experience uniform carriage utilization and maximal theoretical throughput for any queue distribution.

Regardless of the level of control of the queue distributions, a configuration like 251839 is complete for all queue distributions because of the present connectivity between all queues and magazines. The level of control of queues is therefore not important with respect to robustness, but is important in the context of connectivity in order to maintain a specified

$$(SQ) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 0 & 1 \end{pmatrix} \quad (SM) = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \quad (QM) = \begin{pmatrix} 1 & 1 & 2 & 2 \\ 2 & 2 & 3 & 3 \end{pmatrix}$$

Figure 11.15: The connectivity of configuration 251839 results in near maximal throughput for almost all queue distributions because the connectivity is sufficiently high to result in uniform carriage utilization. Additional connections do not significantly improve mean throughput, but do increase the real costs associated with physical constraints.

level of performance.

Control of the queues is also important when considering configurations with absent connectivity between queues and magazines, although the level of control dictates not only the performance but also whether the configuration can transport all items. In general, we can say that control of the queues must increase as the connectivity decreases in order to guarantee completeness and control performance.

Since connectivity is a function of both connections between queues and shafts and connections between magazines and shafts, we can weigh the relative effects of each on robustness and therefore the required level of control of the queue distributions. Identifying how robustness relates to SQ and SM connectivity independently is useful when we consider that practical physical constraints may exist that may make connections between certain system elements more feasible. For instance, in an aircraft weapon elevator system, queues can often have variable locations on the hanger deck, and complete connectivity between all shafts and queues is theoretically possible. However, connectivity between shafts and magazines is typically constrained. Determining how changes in the connectivity between shafts and magazines affect system adaptability and the level of queue control is therefore important for this application.

We have seen how the interaction of shaft-queue and shaft-magazine connectivity collectively affect robustness through the connectivity between queues and magazines. We have seen that the amount of absent connectivity between queues and magazines directly affects the maximum level of robustness for the set of complete evolutions by effectively limiting the number of completely deliverable queue distributions. However, we have not seen any relationship between the *amount* of connectivity for magazines with present connectivity to queues and the achievable level of robustness.

This relationship is presented by summing the number of paths between queues and relevant magazines - those magazines with connectivity to all queues. The connectivity for irrelevant magazines is not included because it can be variable while not affecting completeness and could therefore misrepresent the connectivity of magazines relevant to the set of queue distributions that result in complete evolutions. Representation of the QM connectivity by a single value suffers from some limitations, as compression of the QM connectivity results in the loss of information describing how connectivity is distributed among relevant magazines.

If configurations that halt for at least one queue distribution are omitted because of their misleading effect on the relationship between robustness and connectivity, the level of robustness will only indicate the range of connectivities possible for a given number of relevant magazines, which can still provide a useful characterization of connectivity. Using the set of *unique*, complete evolutions of configurations that do not halt for every queue distribution distributes the set of complete evolutions across various levels of robustness, which illustrates how the number of relevant magazines, and the interaction between the SM and SQ incidence matrices affect the level of robustness. However, we must note that robustness in the context of uniqueness is distinct from robustness in the context of completeness alone. The number of unique evolutions as a measure of robustness describes the "efficiency" of a carriage's presence while the number of complete evolutions describes the ability of a configuration to adapt to various inputs.

Figure 11.16 presents how connectivity between queues and relevant magazines for a configuration is related to the number of unique, complete evolutions the configuration is capable of achieving. The figure indicates that the level of robustness of a configuration, as well as its uniqueness, is limited by the amount of connectivity between its queues and magazines. However, this limitation does not imply that more robust configurations necessarily have high connectivity, only that certain levels of connectivity are required to obtain certain levels of robustness. So, while configuration 262143, with complete connectivity, guarantees that all evolutions will be complete and unique, there exist other configurations with significantly less connectivity, such as configuration 94207, that are still capable of complete, unique evolutions for all queue distributions. Similarly, there exist configurations, such as 258559, that have a relatively large number of connections, yet are unique for a fraction of all queue distributions. The total number of connections alone therefore does not indicate how sparse connectivity leads to maximum robustness or how relatively high connectivity results in relatively low robustness. To see how the connectivity affects robustness, we have to examine the configurations in detail and how the SQ and SM incidence matrices interact.

The incidence matrices for configuration 94207, which corresponds to point A in Figure 11.16 are presented in Figure 11.17 and reveal that the low number of connections between queues and magazines results from sparse connectivity between shafts and queues. The complete connectivity between shafts and magazines guarantees complete evolutions for all queue distributions as well as uniqueness since all carriages must be involved in any evolution. This configuration demonstrates that any configuration with complete connectivity between shafts and magazines results in present connectivity for all magazines and complete, unique evolutions for all queue distributions. Conversely, any configuration without complete connectivity between shafts and magazines is never unique for all queue distributions (although it may be complete for all) since a configuration will not be unique for any queue distribution that is composed solely of item types corresponding to magazines that all lack a connection to at least one common shaft. The number of configurations that are both complete and unique for all queue distributions is therefore a function of the amount of variation possible in the connectivity between shafts and queues under the definitions of validity and uniqueness

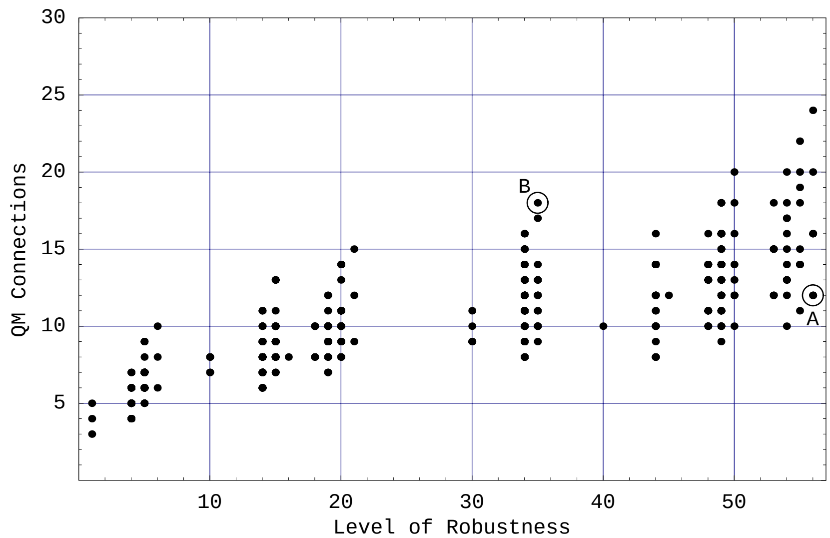

Figure 11.16: The distribution of 2-3-4 configurations with respect to the total number of connections present between all queues and *relevant* magazines and the number of unique, complete evolutions. Analysis of highlighted configurations indicates that the total number alone does not indicate robustness and the actual connections between elements must be considered.

of configurations.

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

Figure 11.17: Despite sparse connectivity between shafts and queues, configuration 94207 is complete and unique for all possible queue distributions because of the complete connectivity between shafts and magazines. Any configuration with complete connectivity between shafts and magazines results in complete, unique evolutions for any queue distribution, although throughputs may vary depending on the SQ connectivity. The number of configurations with complete SM connectivity is dependent on the valid variations possible for connectivity between shafts and queues.

In contrast to configuration 94207 is configuration 258559, which corresponds to point B in Figure 11.16 and has the connectivity illustrated in Figure 11.18. The large number of connections between queues and shafts is largely a function of the complete connectivity between shafts and queues. Despite this connectivity, only 35 queue distributions are unique for 258559 because any queue distribution lacking item types corresponding to the fourth magazine results in no utilization of the first carriage. The sparse connectivity of the first shaft also indicates that throughput is relatively low for the 21 queue distributions that lack an item type corresponding to the fourth magazine because only two carriages are utilized. And, in fact, configuration 258559 represents one of the low mean throughput/high connectivity, or "inefficient" configurations in Figure 11.11.

$$(SQ) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 1 & 1 \end{pmatrix} \quad (SM) = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \quad (QM) = \begin{pmatrix} 2 & 2 & 2 & 3 \\ 2 & 2 & 2 & 3 \end{pmatrix}$$

Figure 11.18: Configuration 258559 has complete connectivity between shafts and queues, which results in a large number of connections between queues and magazines. However, the sparse connectivity of the first shaft with respect to magazines means that any queue distribution lacking items bound for the fourth magazine will result in a non-unique evolution. The sparse connectivity also means that throughput for these queue distributions will be low, because only two carriages are utilized.

Although we have been considering only those magazines with present connectivity to all queues in determining the number of connections between queues and magazines, accounting for the connectivity of all magazines has negligible impact on the fundamental relationship between connectivity and robustness evident in Figure 11.16. Figure 11.19 presents the distribution of configurations with respect to robustness and connectivity for all magazines, relevant and irrelevant. At higher levels of robustness (22 and greater), there are, by definition, no differences because the number of relevant magazines is equal to the number of magazines in the system. At lower levels of robustness, where configurations with irrelevant magazines dominate, the connectivity increases, depending on the number of irrelevant magazines. However, the differences are not significant enough to result in any qualitative differences in the relationships in Figure 11.16 and 11.19. The lack of any significant differences between connectivities for configurations with irrelevant magazines suggests that the connectivity of irrelevant magazines follows the connectivity of relevant magazines, a condition not necessary for irrelevance.

Even with a common basis for comparisons of the connectivity among configurations with different numbers of relevant magazines, comparisons of configurations with a common number of relevant magazines are inherently more valid because present connectivity is intimately related to the set of possible queue distributions that can result in complete and unique evolutions and therefore robustness. So, when we compare configurations with different numbers of relevant magazines, we consider the same number of magazines, but consider different sets of queue distributions.

Figure 11.20 presents the distribution of configurations that are unique and complete for various numbers of queue distributions and have relevant connectivity for all magazines. The figure shows that robustness is still limited by the amount of connectivity. However, configurations with relatively low connectivity are still capable of achieving relatively high levels of robustness. It is therefore possible to indicate the minimum possible level of robustness for a given amount of connectivity and number of relevant magazines. Figure 11.20, when compared with either Figure 11.16 or Figure 11.19 also reveals a form of nesting that is a

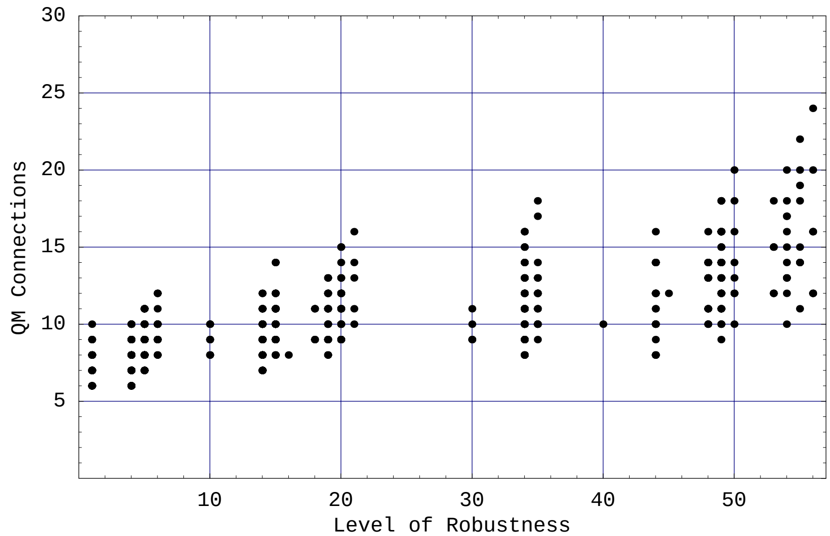

Figure 11.19: The distribution of 2-3-4 configurations with respect to the number of complete, unique evolutions and the total number of connections between queues and all magazines, relevant and irrelevant. The relationship is fundamentally identical to that for the distribution when only the connections for relevant magazines are considered, suggesting that the connectivity of irrelevant magazines follows the connectivity of relevant magazines.

function of the amount of present connectivity and the effects of the interactions between the SQ and SM incidence matrices on robustness.

For any set of configurations with a common number of relevant magazines, the shape of the relationships between the number of connections between queues and magazines and robustness are quite similar. This similarity is evident in a comparison of Figure 11.20 with Figure 11.21, which presents the relationships between robustness and the amount of connectivity between queues and magazines for *(a)* configurations with three relevant magazines and *(b)* configurations with two relevant magazines.

The analysis of configuration 94207 (point A in Figure 11.16) revealed that complete connectivity between shafts and magazines guaranteed complete and unique evolutions, regardless of the queue distribution. The connectivity between shafts and queues can therefore be any possible combination, as long as the connectivity is valid and does not result in a redundant configuration. As the connectivity between shafts and magazines decreases however, the connectivity between shafts and queues becomes more limited in order to maintain present connectivity for all magazines. When the connectivity between shafts and magazines becomes sufficiently low, only complete connectivity between shafts and queues results in present connectivity for all magazines. Since only complete SQ connectivity is possible, the SM connectivity is the minimum possible for any configuration and any further decrease results in an invalid configuration. Complete SQ connectivity always results in present connectivity to queues for all magazines, regardless of the SM connectivity. Complete SQ or SM

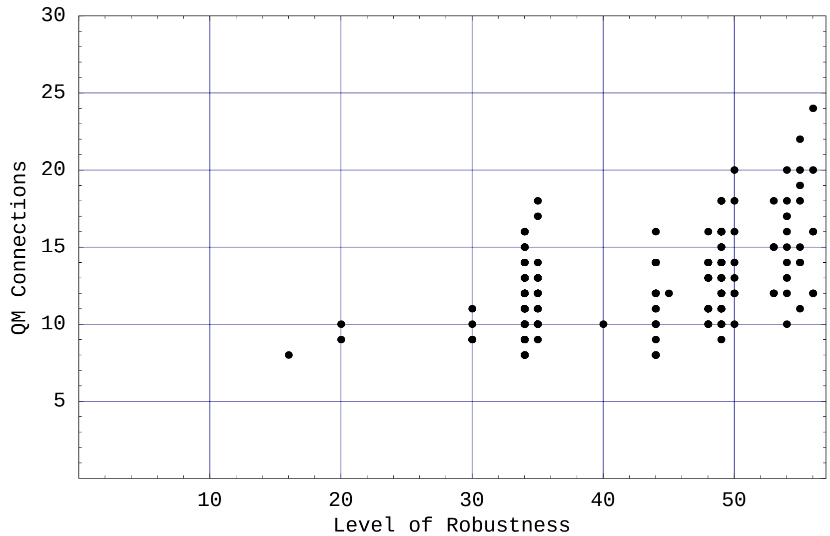

Figure 11.20: The distribution of 2-3-4 configurations that have present connectivity for all magazines with respect to robustness and the number of connections between queues and magazines. All configurations with present connectivity for the same number of magazines are comparable for the same set of queue distributions. Separating configurations according to the number of magazines with present connectivity also illustrates a nesting behavior between connectivity and robustness.

connectivity therefore guarantee complete evolutions for all queue distributions, but only complete SM connectivity guarantees uniqueness as well.

For configurations with fewer magazines with present connectivity to all queues, the relationship between SM and SQ connectivity repeats, except the robustness levels are dependent on the number of relevant magazines. For example, with three relevant magazines, a maximum of 21 of our queue distributions can result in complete and unique evolutions. At this maximum level of robustness, only configurations with complete connectivity between shafts and *relevant* magazines are possible. Again, the number of configurations with complete connectivity for a number of relevant magazines is dependent on the number of valid combinations of connectivity between shafts and queues, ranging from near complete, but sufficiently low to result in less than the maximum number of relevant magazines, to minimal. The number of configurations with each variation of SQ connectivity is in turn dependent on the connectivity between shafts and *irrelevant* magazines. As the connectivity between shafts and relevant magazines decreases, the connectivity between shafts and queues must increase in order to maintain present connectivity to all queues for all relevant magazines. However, SQ connectivity between shafts and queues tends to decrease as the number of relevant magazines decreases because greater SQ connectivity increases the potential for present connectivity between queues and magazines.

The effects of the relationship between connectivity of shafts and magazines and the connec-

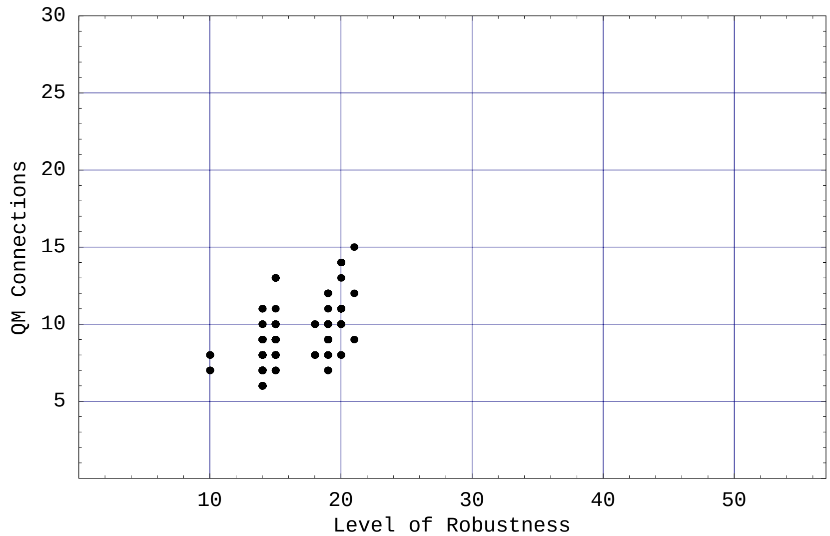

(a)

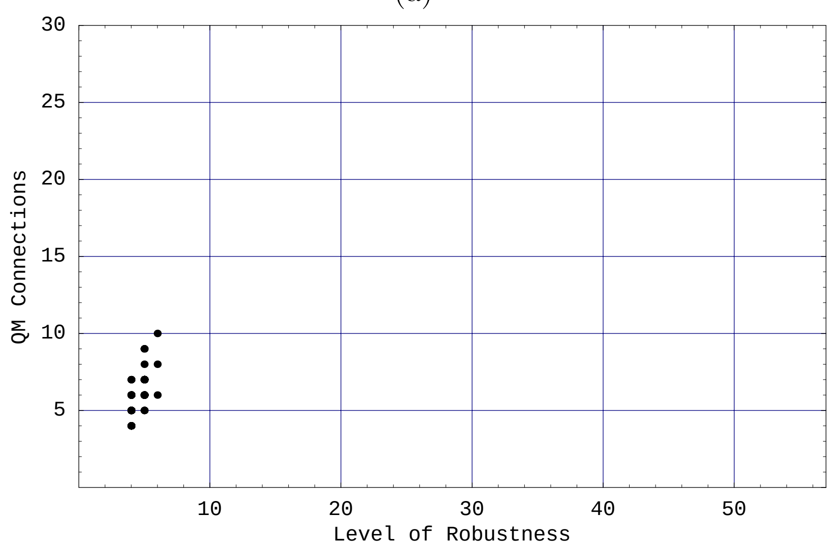

(b)

Figure 11.21: The distribution of 2-3-4 configurations with respect to robustness and QM connectivity when (a) only three magazines have present connectivity and (b) only two magazines have present connectivity. The distributions have a similar shape yet exist at different ranges of robustness, resulting in a nested distribution when all distributions of configurations with different numbers of relevant magazines are superposed.

tivity of shafts and queues on robustness, with respect to the number of relevant magazines, is evident when we examine the ranges of connectivity in both the SQ and SM incidence matrices of configurations at various levels of robustness. Figure 11.22 presents the number of connections between shafts and relevant magazines and the average number of connections between shafts and queues for all configurations that result in complete, unique evolutions for various numbers of queue distributions. The average number of connections between shafts and queues for all configurations at a level of robustness is used rather than the range of values because configurations with sparse SM and SQ connectivity are possible, particularly with fewer relevant magazines. With no indication of the frequency of connectivities, the relationship between relevant magazines, robustness, and SQ connectivity is obscured. While the relative differences in the amount of SM and SQ connectivities vary at different levels of robustness, the figure illustrates that greater SM connectivity occurs at the maximum levels of robustness for various numbers of relevant magazines (1^{st}, 6^{th}, 21^{st}, and 56^{th}) and that the average SQ connectivity must increase between these values to maintain a number of relevant magazines.

Since the minimum SM connectivity for configurations with a particular number of relevant magazines may result in fewer complete, unique evolutions than a configuration with fewer relevant magazines, examining sets of configurations with a common number of relevant magazines presents a clearer relationship between SQ and SM connectivity and robustness. Figure 11.23 shows the number of connections in the SM (relevant magazines) matrices and the average number of connections in the SQ incidence matrices at different levels of robustness for configurations with various numbers of relevant magazines. For any set of configurations with a common number of relevant magazines, the greatest level of robustness corresponds to configurations that have complete connectivity for relevant magazines and a minimum or near minimum average connectivity between shafts and queues. Decreasing the connectivity between shafts and relevant magazines decreases robustness and increases the average connectivity between shafts and queues, although not necessarily in a regular way - some amount of nesting is apparent for certain sets of configurations.

However, the relationship between average connectivities for the SQ and SM incidence matrices, irrespective of robustness, is apparent for any number of relevant magazines. Isolating configurations with common numbers of relevant magazines presents a clear picture of this relationship, evident in Figure 11.24 for configurations with four, three, and two relevant magazines. The inverse relationship results from the necessity for more complete SQ connectivity as SM connectivity decreases for a number of relevant magazines in order to maintain the number of relevant magazines.

The nesting apparent in Figure 11.23(a) (and to a lesser extent in (c)), which is not apparent in a direct comparison of average connectivities of SQ and SM incidence matrices at each level of robustness, is a result of the combined influence of the number of connections and the details of how connections are distributed, information that is not representable by a single value. As mentioned previously, a configuration with complete connectivity between shafts and magazines results in unique evolutions for all queue distributions. Lower levels

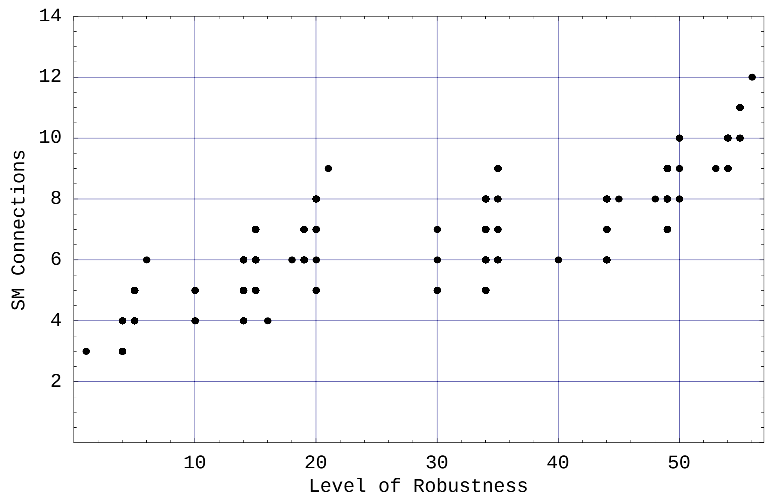

(a)

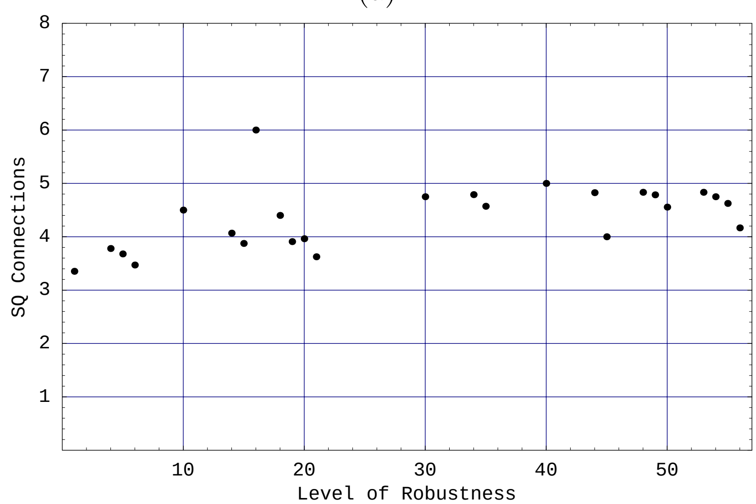

(b)

Figure 11.22: The (a) number of connections for SM incidence matrices and (b) the average number of connections for SQ incidence matrices for 2-3-4 configurations at various levels of robustness. The number of SM connections accounts for the connectivity of relevant magazines. As the number of connections between shafts and magazines falls, the number of connections between shafts and queues must increase in order to obtain a certain number of relevant magazines. The average SM connectivity is maximum for configurations with a number of relevant magazines at the most robust level associated with that number of magazines.

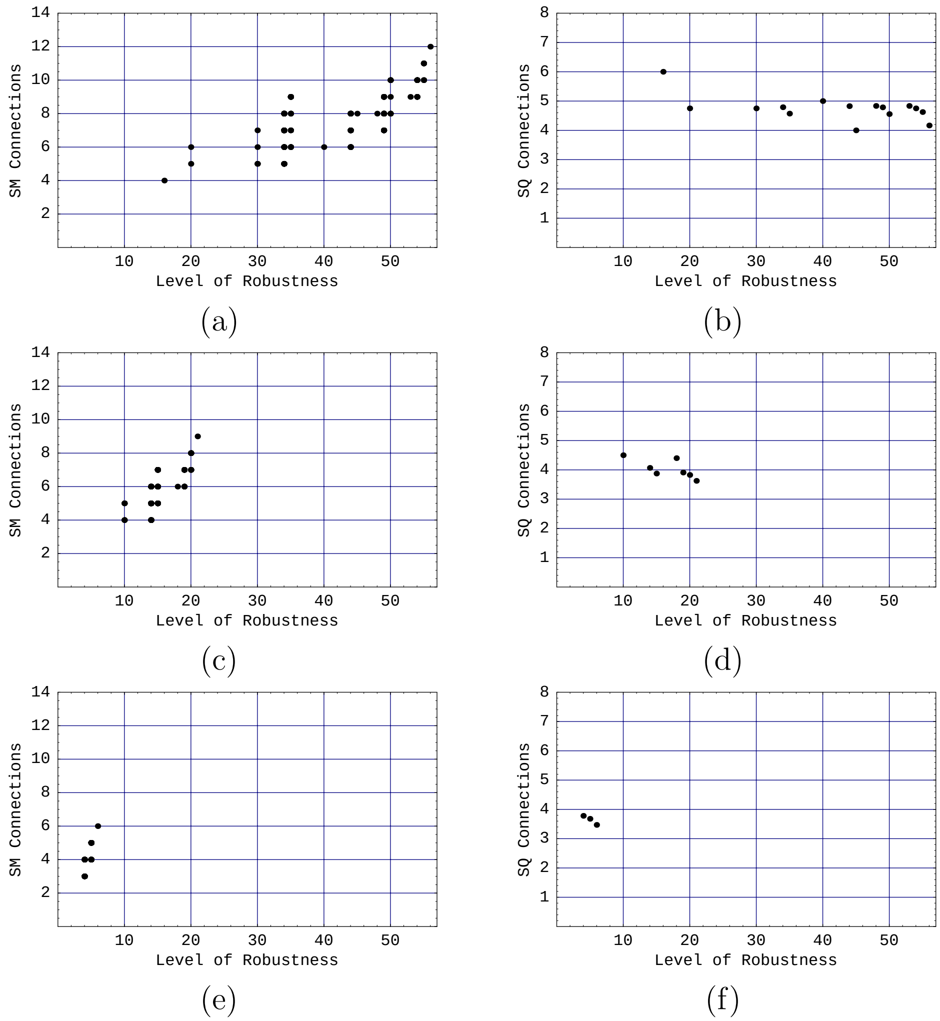

Figure 11.23: Relationships between the number of connections in the SM incidence matrices and robustness and between the average number of connections in the SQ matrices and robustness for 2-3-4 configurations with common numbers of relevant magazines. Each successive row corresponds to the distributions for configurations with a common number of relevant magazines, starting with four and decreasing to two. Isolation of the configurations with a common number of relevant magazines provides a clearer illustration of the interaction between SM and SQ connectivities.

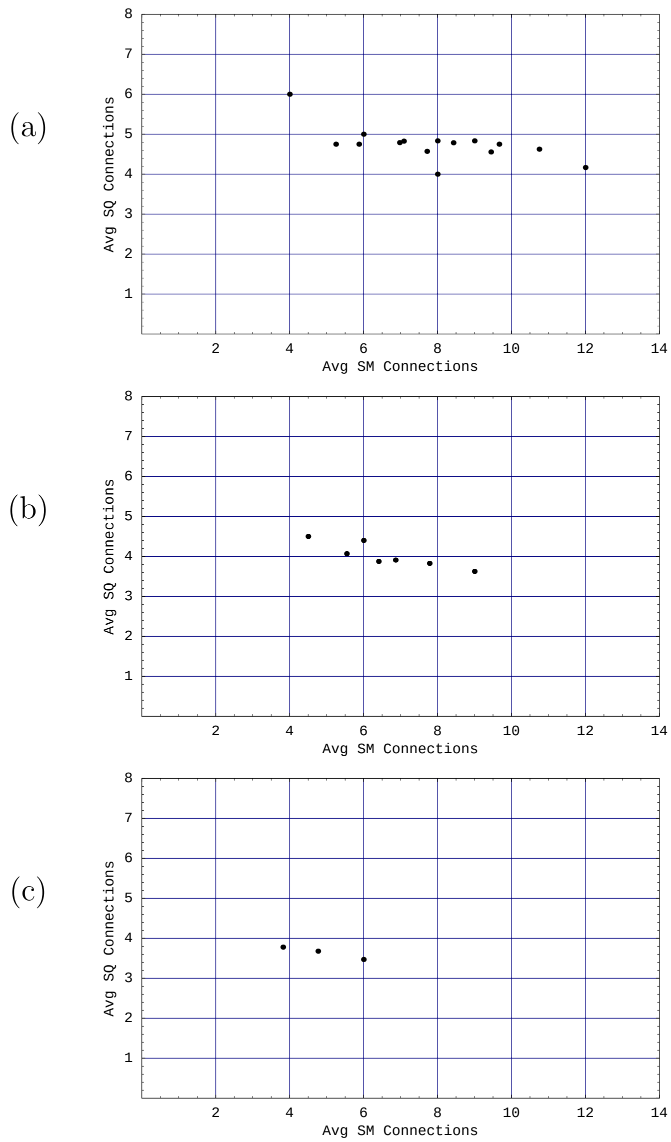

Figure 11.24: Comparisons of average numbers of connections in the SQ and SM incidence matrices irrespective of robustness for 2-3-4 configurations with (a) four relevant magazines, (b) three relevant magazines, and (c) two relevant magazines. The figures show that the SM and SQ connectivities are inversely related because lower SQ connectivities can only occur when SM connectivities are sufficiently high in order to maintain a specified number of relevant magazines.

of robustness are a function of the number of magazines that have absent connections to shafts and the number of absent connections for each shaft. If a single magazine lacks a connection to one or more, but not all shafts, then the queue distribution consisting entirely of items corresponding to the magazine results in a non-unique evolution. If a shaft has absent connectivity to multiple magazines, then any queue distribution consisting entirely of combinations of item types bound for these unconnected magazines results in non-unique evolutions, including those composed of a single item type.

Figure 11.25 presents the connectivities of configurations 259967 and 260029 (both 2-3-4 size systems), which correspond to the 55^{th} and 53^{rd} levels of robustness, respectively. The SQ incidence matrices are complete in both cases to facilitate comparison. For configuration 259967, the first magazine has absent connectivity shafts. As a result, the queue distribution with all items bound for the first magazine (100-0-0-0) leads to a non-unique evolution because the first and second magazines are not utilized. For configuration 260029, each magazine has a single absent connection to a distinct shaft. As a result, three queue distributions, each composed entirely of items bound for either of the magazines with an absent connection result in non-unique evolutions. Note that robustness is unaffected by the number of shafts that are not connected to a single magazine, as long as additional absences do not occur for a shaft with another absent connection to another magazine. Configurations with different numbers of connections can therefore share the same level of robustness. So, configurations 259967 could have present connectivity between the second shaft and the first magazine (which is configuration 260095) and still result in the same level of robustness. However, configuration 260029 can not have any additional absent connections for any magazine because an additional absence anywhere would result in a shaft with multiple absent connections, which decreases robustness. The limit on additional absences places a lower limit on the levels of robustness for configurations that result in non-unique evolutions only for queue distributions composed entirely of a single item type. This limit based on the lower value of either the number of shafts or magazines. Therefore, for our example 2-3-4 system, in which there are fewer shafts than magazines, the minimum level of robustness possible for configurations that are non-unique only for queue distributions composed of a single item type is $56 - 3 = 53$.

Figure 11.26 presents the connectivity of configuration 259071, in which the first shaft lacks a connection to both the first and second magazines. This configuration therefore results in non-unique evolutions for any queue distribution that is composed of any combination of items corresponding only to these magazines, including queue distributions composed entirely of a single item type corresponding to either of these magazines. There are six combinations of distributions involving these item types (the same six that may result in robust evolutions with present connectivity for two magazines) and the configuration therefore exists at the $56 - 6 = 50^{th}$ level of robustness.

Although configuration 259071 has a lower level of robustness than configurations that are non-unique only for queue distributions that consist entirely of a single item type, the number of connections in configuration 259071 can be greater simply because of the greater number

$$(SQ) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 1 & 1 \end{pmatrix} \quad (SM) = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \quad (QM) = \begin{pmatrix} 1 & 3 & 3 & 3 \\ 1 & 3 & 3 & 3 \end{pmatrix}$$

(a)

$$(SQ) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 1 & 1 \end{pmatrix} \quad (SM) = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \end{pmatrix} \quad (QM) = \begin{pmatrix} 2 & 2 & 2 & 3 \\ 2 & 2 & 2 & 3 \end{pmatrix}$$

(b)

Figure 11.25: Incidence matrices for configurations (a) 259967 and (b) 260029, both of which are of the 2-3-4 system size. When magazines have absent connections to at least one shaft, and the corresponding shaft(s) are unconnected to only one magazine, then only the queue distributions composed entirely of item types corresponding to these magazines result in non-unique evolutions.

$$(SQ) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 1 & 1 \end{pmatrix} \quad (SM) = \begin{pmatrix} 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \quad (QM) = \begin{pmatrix} 2 & 2 & 3 & 3 \\ 2 & 2 & 3 & 3 \end{pmatrix}$$

Figure 11.26: The incidence matrices for configuration 259071 of system size 2-3-4. The first shaft has absent connections to two magazines, so queue distributions composed of any combination of item types corresponding only to these magazines result in non-unique evolutions. Configurations with fewer connections can have greater robustness when no shaft has multiple absent connections, which explains the nesting in Figure 11.23

of queue distributions consisting of combinations of multiple item types that exist. Similarly, configuration 258559, with connectivity presented in Figure 11.27, is not unique for the set of queue distributions formed by combinations of only three item types, of which there are 21. We studied configuration 258559 earlier with respect to QM connectivity and throughput, recognizing it as a low mean throughput/high connectivity, or "inefficient" configuration. While the number of connections between shafts and magazines in configuration 258559 is only one less than that for 259071, the configuration results in unique evolutions for only $56 - 21 = 35$ queue distributions, 35 less than that for configuration 259071. Between these levels of robustness, configurations run through all combinations of connectivity that are possible given that no shaft is not connected to more than two magazines.

$$(SQ) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 1 & 1 \end{pmatrix} \qquad (SM) = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \qquad (QM) = \begin{pmatrix} 2 & 2 & 2 & 3 \\ 2 & 2 & 2 & 3 \end{pmatrix}$$

Figure 11.27: Further evidence of the reasons for nesting is the connectivity of configuration 258559, 2-3-4. Compared to configuration 259071, configuration 258559 has only one additional absent connection. However, because a shaft has absent connectivity to three magazines, significantly fewer queue distributions result in unique evolutions.

The maximum number of absent connections to magazines for any shaft and the connectivity of remaining shafts therefore creates the nesting observed in the relationship between robustness and connectivity in SM incidence matrices for a set of configurations with a common number of relevant magazines. On top of this nested relationship is the nesting associated with the number of relevant magazines when all configurations with unique, complete evolutions are considered.

Identifying the causes behind nesting in the distribution of configurations with respect to uniqueness illustrates how the actual connections, not just the number of connections are important not only for determining uniqueness, but for completeness, which is arguably a better description of robustness. Two configurations may share the identical number of connections, but the actual connections can result in different numbers of relevant magazines and therefore significantly different numbers of complete evolutions. But since we recognize that robustness is a function of the number of relevant magazines, we can use connectivity as a configuration filter - only considering those configurations with present connectivity between all queues and magazines. We have also seen how the actual connections and not just the number of connections must be considered with respect to throughput, even when configurations have the same number of relevant magazines. A configuration with redundant shafts, which have identical connectivity to both queues and magazines, will have lower mean throughput for a variety of queue distributions when compared to a configuration with the same number of connections, but no redundancy. Redundancy can therefore also be applied as a filter for identifying potentially optimal configurations.

The causes of nesting also provide insight on how the SQ and SM incidence matrices interact to result in a particular number of relevant magazines and therefore robustness. If constraints exist on the connectivity between shafts and magazines, then there can exist limitations on the minimum amount of connectivity between shafts and queues in order to have a particular number of relevant magazines and a certain level of robustness. Constraints on connectivity must also be considered with respect to throughput. Figure 11.11 shows that limitations exist on the maximum mean throughput, given a particular number of possible connections. Attaining the maximum is dependent on the actual connections of the configuration. However, near maximal mean throughputs are achievable by configurations with less than complete connectivity that exist at the “knee” in Figure 11.11. Depending on the cost of connectivity, configurations in this region are optimal with respect to performance.

The principal advantage of using connectivity in the search for optimal configurations is that it is a static measure and explicit simulation is not required. Furthermore, based on what we have seen in the analysis of connectivity, it is possible to build a configuration to suit constraints such as acceptable performance, limitations on physical connections, the incremental cost of connections, and the level of control of queues. However, expression of connectivity as a single value (average connectivity, number of connections) alone provides only an indication of potential performance and robustness because too much information regarding the actual arrangement of a configuration is lost in the compression. To fully utilize connectivity, additional detailed measures are required that increase the costs associated with optimization.