ELectronics EXpress

IEICE Electronics Express, Vol.4, No.21, 657–664

Voting structures for cascaded triple modular redundant modules

Sungjae Lee, Jae-il Jung, and Inhwan Lee[a)]

Department of Electrical and Computer Engineering, Hanyang University, Hangdang-Dong 17, Sungdong-Gu, Seoul, 133–791, Korea

a) *ihlee@hanyang.ac.kr*

Abstract: This paper investigates what kinds of choices exist in determining the voting structure for cascaded triple modular redundant (TMR) modules, how to identify efficient voting structures, and the effects of voting structures on the overall system reliability. While the classic single-voter and three-voter architectures have been used for about fifty years, the paper shows that there are more practically useful voting structures which provide efficient trade-offs between hardware overhead and system reliability. Specifically, a single-voter architecture with the voter fanout of one is found to be more reliable than the classic single-voter architecture for a reasonably wide range of component reliability.

Keywords: cascaded triple modular redundant (TMR) modules, voting structure, hardware overhead, reliability, recovery distance

Classification: Science and engineering for electronics

References

[1] J. V. Neumann, "Probabilistic logics and the synthesis of reliable organisms from unreliable components," in *Automata Studies*, ed. C. E. Shannon and J. McCarthy, pp. 43–98, Princeton University Press, Princeton, 1956.

[2] W. G. Brown, J. Tierney, and R. Wasserman, "Improvement of electronic-computer reliability through the use of redundancy," *IRE Trans. Elec. Comp.*, vol. EC-10, no. 3, pp. 407–416, Sept. 1961.

[3] K. J. Gurzi, "Estimates for best placement of voters in a triplicated logic network," *IEEE Trans. Electron. Comput.*, vol. EC-14, no. 5, pp. 711–717, Oct. 1965.

[4] S. R. McConnel and D. P. Siewiorek, "Synchronization and voting," *IEEE Trans. Comput.*, vol. C-30, no. 2, pp. 161–164, Feb. 1981.

[5] D. P. Siewiorek and R. S. Swarz, *The Theory and Practice of Reliable System Design*, Digital Press, Bedford, MA, 1982.

[6] G. Latif-Shabgahi, J. M. Bass, and S. Bennett, "A taxonomy for software voting algorithms used in safety-critical systems," *IEEE Trans. Reliab.*, vol. 53, no. 3, pp. 319–328, Sept. 2004.

[7] S. Lee and I. Lee, "Staggered voting for TMR shift register chains in poly-Si TFT-LCDs," *Journal of Information Display*, vol. 2, no. 2, pp. 22–26,

DOI: 10.1587/elex.4.657
Received September 17, 2007
Accepted September 28, 2007
Published November 10, 2007

IEICE Electronics Express, Vol.4, No.21, 657–664

June 2001.

1 Introduction

The concept of triple modular redundant (TMR) architecture has been proposed long time ago as a method to build reliable systems using unreliable components [1, 2, 3]. Benefits of implementing a complex system as a network of TMR modules, by partitioning the system into smaller subsystems, have also been recognized. Two classic voting structures - the single-voter and the three-voter architectures - have been established. Detailed issues in synchronization and voting have been investigated [4]. Many voting methods have been developed and used for building real high-reliability systems [5, 6].

This paper revisits the voting structure for cascaded TMR modules. In this work, a TMR stage (or a stage for short) refers to a triplicated module (M) and the associated voters (V), and cascaded TMR modules refers to a series connection of multiple TMR stages (Fig. 1). Cascaded TMR modules can constitute a complete system. Or it can be a subsystem of a larger, more general triplicated logic network. That is, the output of a stage in Fig. 1 can be used as inputs of other subsystems, and the outputs of other subsystems can also be used as additional inputs of some stages in Fig. 1. Cascaded TMR modules are considered to be reliable if at least two of the three module inputs are correct in every stage and the last-stage voter generates correct output.

The two classic voting structures in Figs. 1 (a) and 1 (b) have been used for about fifty years. In these architectures, every stage, except possibly the last one, has the same structure - we assume that the last stage always uses a single voter. The single-voter architecture uses one majority voter (or voter for short) per stage. Since a majority voter can generate correct output if two out of three inputs are correct, the single-voter architecture can tolerate a module failure. But it cannot tolerate a voter failure. The three-voter architecture uses more voters and can tolerate a voter failure as well.

However, our previous study indicated that there is another practically useful voting structure for cascaded TMR modules [7]. Specifically, the study used the structure in Fig. 1 (c) for building a long chain of TMR shift registers for poly-silicon TFT-LCD (thin-film transistor liquid crystal display) driving. The rationale for this voting structure is that, while it uses a single voter per stage, a voter failure does not immediately cause a system failure because a voter affects only one module. This is an important advantage over the classic single-voter architecture. A voter failure can be masked if there are no additional failures in the next three stages. Similarly, a module failure can be masked within the next two stages if there are no additional failures.

The study showed that Fig. 1 (c) is almost as reliable as the classic three-voter architecture when the module and the voter are equally reliable and there are several hundreds of stages in the system. This result indicates that there is more freedom in determining the voting structure for cascaded TMR

DOI: 10.1587/elex.4.657
Received September 17, 2007
Accepted September 28, 2007
Published November 10, 2007

ELectronics EXpress

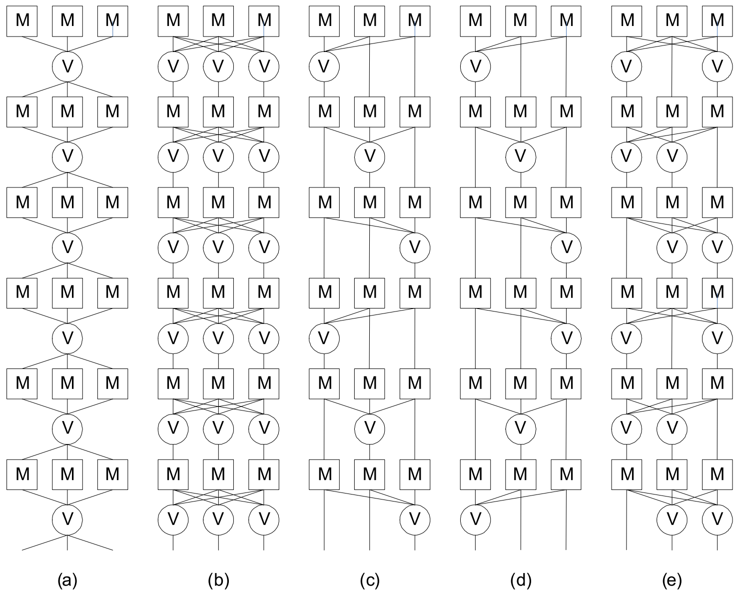

Fig. 1. Voting structures for cascaded TMR modules: (a) classic single-voter, (b) classic three-voter, (c) single-voter with voter fanout of one, (d) another single-voter with voter fanout of one, (e) two-voter with voter fanout of one

modules. This paper expands and generalizes the work. Specifically, the paper investigates what kinds of freedom exist, how to find efficient voting structures, and the effects of voting structures on the overall system reliability. The paper also investigates what happens when the voter is more reliable than the module and when there are only a few stages in the system.

2 Voting structures

While the classic single-voter and three-voter architectures have been well-established, we have more choices in determining the voting structure for cascaded TMR modules. The first choice we have is the number of voters per stage. If we use less than three voters per stage, we also need to determine the voter fanout. The term voter fanout refers to the number of modules which receive the output of a voter. Additionally, we need to determine the fanout topology. That is, we need to determine how to connect the voter outputs from a stage with the modules in the next stage. On the other hand, the voting logic is always two-out-of-three, and the voter always uses the outputs of all three modules.

Let's first consider the use of one voter per stage. If the voter fanout is three, it becomes the classic single-voter architecture in Fig. 1 (a). If the voter fanout is one, two modules in a stage directly receive the module outputs from the previous stage, without voting, as in Fig. 1 (c) – we will discuss more about this case below. If the voter fanout is two, a voter failure means a system failure because a voter failure affects two modules in the next stage. Also, the effects of module failures may propagate to the next stage because one module in a stage directly receives the module output from the previous

DOI: 10.1587/elex.4.657
Received September 17, 2007
Accepted September 28, 2007
Published November 10, 2007

IEICE Electronics Express, Vol.4, No.21, 657–664

stage. Therefore, such a voting structure is not better than the classic single-voter architecture.

When the voter fanout is one, there are many possible fanout topologies. However, it is not straight-forward to determine if one fanout topology is better than another. To facilitate this comparison, we introduce the notion of *recovery distance for voter failure* (RDVF). It is defined as the number of stages required to recover from a voter failure. RDVF is generally represented by a sequence of numbers. But it is always three in Fig. 1 (c) because the effect of a failed voter can be masked by the next voter in the same column and the distance between two successive voters in a column is always three. In a sense, RDVF represents the size of vulnerable window. Given a voter failure in Fig. 1 (c), if additional failures occur in the next three stages, a system failure may occur. Otherwise, the voter failure is masked. Thus, the size of vulnerable window is three stages. RDVF can be a useful measure for comparing fanout topologies.

We can also consider the *recovery distance for module failure* (RDMF), which is the number of stages required to recover from a module failure. Let's consider the failure of a module which immediately follows a voter in Fig. 1 (c). Then, if additional failures occur in the same stage or in the next two stages, a system failure may occur. Otherwise, the module failure is masked. Thus, the size of vulnerable window is three stages. For some modules, depending on their position with respect to the neighboring voters, the size of vulnerable window can be either two stages or one stage. Specifically, the RDMF for Fig. 1 (c) is represented by the sequence of 3, 2, and 1. Note that, since a module failure has to be recovered eventually by a voter in Fig. 1 (c), the RDMF of a module is the distance between the module and the first voter which follows the module in the same column. Thus, given a module, its recovery distance is bounded by the RDVF of the voter which immediately precedes the module in the same column.

Let's investigate the notion of recovery distance further using a different fanout topology in Fig. 1 (d). It uses the same number of voters with Fig. 1 (c). In Fig. 1 (d), the structure of the ith stage is identical to that of the $(i{+}6)$th stage. The RDVF for the first and third columns is represented by the sequence of 5 and 1. The RDMF for the first and third columns is represented by the sequence of 1, 5, 4, 3, 2, and 1. The recovery distance characteristics for the second column are identical to those in Fig. 1 (c).

In both Figs. 1 (c) and 1 (d), the average RDVF is 3.0, which is always true when we use one voter per stage. But the maximum value of RDVF is 5 in Fig. 1 (d), which causes larger values of RDMF such as 4 and 5. Thus, the average RDMF for all columns in Fig. 1 (d) is 2.44, while it is 2.0 in Fig. 1 (c). In fact, Fig. 1 (c) is one of the architectures which have the shortest average RDMF - note that, by changing the order of stages in Fig. 1 (c), we can generate more fanout topologies which have the same recovery distance characteristics with Fig. 1 (c). A shorter recovery distance indicates that system failures can be less likely to occur. The simulation results confirm that Fig. 1 (c) is more reliable than Fig. 1 (d). The results of

DOI: 10.1587/elex.4.657
Received September 17, 2007
Accepted September 28, 2007
Published November 10, 2007

IEICE Electronics Express, Vol.4, No.21, 657–664

simulation of a variety of fanout topologies show that the average RDMF is a useful measure for comparing fanout topologies.

We rely on RDVF and RDMF to compare different fanout topologies. That is, we select a smaller set of candidates from all possible fanout topologies based on the maximum values of RDVF and RDMF. Then we examine the average RDMF of each candidate to finalize the fanout topology. We do not expect a random fanout topology to be useful because the maximum value of RDVF for a random topology will be larger than that for a regular topology such as Fig. 1 (c).

Let's consider the use of two voters per stage. If the voter fanout is one, we can show that Fig. 1 (e) is one of the architectures which have the best recovery distance characteristics. Its RDVF is represented by the sequence of 1 and 2, and its RDMF by the sequence of 1, 1, and 2. As in Fig. 1 (c), the structure of ith stage is identical to that of $(i{+}3)$th stage in Fig. 1 (e). The architectures in which the fanout of one voter is one and the fanout of the other voter is two are not better than the classic single-voter architecture. Such architectures cannot tolerate the failure of the voter with two fanouts, yet allow the failure of the voter with one fanout to propagate to the next stage.

Cascaded TMR modules with three voters per stage are the classic three-voter architecture in Fig. 1 (b). If we use more than three voters per stage, the additional voters can only be used as spares.

3 Reliability evaluation

Cascaded TMR modules are considered to be reliable if at least two of the three module inputs are correct in every stage and the last-stage voter generates correct output. The reliability of the two classic architectures in Figs. 1 (a) and 1 (b) can be evaluated analytically [5]. However, we cannot readily do that for the rest of the architectures because it can take multiple stages to recover from a component (voter or module) failure in those architectures. That is, there are inter-stage failure dependencies. We performed a probabilistic calculation in our previous work [7], but realized later that such a calculation does not capture all failure dependencies. So we use the Monte Carlo simulation to perform the evaluation.

Let's illustrate the simulation process using the architecture in Fig. 1 (c). Each simulation run starts from the first stage. Given a stage, if two or more modules fail, then the stage fails. If both a voter and a module in different columns fail, the stage also fails. If a stage does not fail, the effect of any failure in the current stage, if exists, is propagated to a corresponding module in the next stage, and we move to the next stage. Here, a stage failure means the failure of the entire system. Once a stage failure occurs, the current simulation run is terminated and the next run is started.

To facilitate reliability evaluation, we assume that all modules have the same reliability and all voters have the same reliability. We also assume that the last stage always uses a single voter.

DOI: 10.1587/elex.4.657
Received September 17, 2007
Accepted September 28, 2007
Published November 10, 2007

IEICE Electronics Express, Vol.4, No.21, 657–664

Fig. 2 shows the unreliability of Figs. 1 (a), 1 (b), 1 (c), and 1 (e) for several values of $(1\text{-}R_m)/(1\text{-}R_v)$. The term unreliability is defined as "1-reliability." R_m and R_v are the module reliability and the voter reliability, respectively. Thus, $(1\text{-}R_m)/(1\text{-}R_v)$ is the ratio of module unreliability to voter unreliability. As its value increases, a module becomes more likely to fail than a voter. In the simulation, the value of R_v is fixed at $1\text{-}10^{-5}$and we perform simulation for various values of R_m. The same trend is observed for different values of R_v. Generally, the output of a module can be a binary value, an analog value, or an arbitrary sequence of values. Also, the voter may have to synchronize the module outputs before voting. As a result, a voter can be as simple as a few logic gates or it can be a complex processor-based system. The number of stages in cascaded TMR modules is assumed to be 100. Fig. 2 also shows the unreliability of non-redundant cascaded modules.

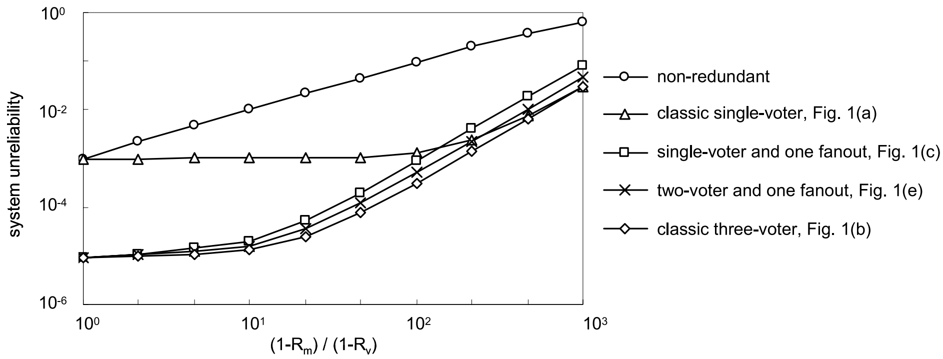

Fig. 2. Reliability of cascaded TMR modules

As reported in [7], when R_m and R_v are identical, Fig. 1 (c) uses one voter per stage but is almost as reliable as the classic three-voter architecture. A voter failure does not immediately cause a system failure in Fig. 1 (c), which is why Fig. 1 (c) is much better than the classic single-voter architecture. When compared with the classic three-voter architecture, module failures may propagate to the next stage and thus two component failures in the neighboring stages are more likely to cause a system failure in Fig. 1 (c). However, Fig. 1 (c) uses a fewer number of voters, so fewer voter failures will occur. The end result is that Fig. 1 (c) is almost as reliable as the classic three-voter architecture. The data shows that, when both the values of R_m and R_v are $1\text{-}10^{-5}$ (that is, at the left end of Fig. 2), system failures due to double component failures are rare in Figs. 1 (c), 1 (e), and 1 (b), and their reliabilities are determined mainly by the reliability of the last-stage voter. As a result, using more voters does not help. The classic single-voter architecture cannot tolerate a voter failure, so it is at best as reliable as non-redundant architecture when R_m and R_v are identical.

As a module becomes more likely to fail than a voter, the reliability gaps among Figs. 1 (c), 1 (e), and 1 (b) become wider, indicating that using more voters helps. Still, Figs. 1 (c) and 1 (e) remain to be more reliable than the

DOI: 10.1587/elex.4.657
Received September 17, 2007
Accepted September 28, 2007
Published November 10, 2007

classic single-voter architecture until a module becomes over a hundred times more likely to fail than a voter. Clearly, they are practically useful architectures which can provide efficient trade-offs between hardware overhead and system reliability. If the value of $(1\text{-}R_m)/(1\text{-}R_v)$ increases further, Figs. 1 (c) and 1 (e) become less reliable than the classic single-voter architecture.

To investigate this reliability cross, we analyze the failure scenarios using the data collected from the Monte Carlo simulation. The results show that, when a module is much more likely to fail than a voter, propagating the output of a module to the next stage, without passing a voter, significantly increases the chance of double module failures in a stage, i.e., the chance of system failure. That is why Figs. 1 (c) and 1 (e) become less reliable than the classic single-voter architecture. When double module failures are common, the classic single-voter and three-voter architectures are not much different in terms of reliability.

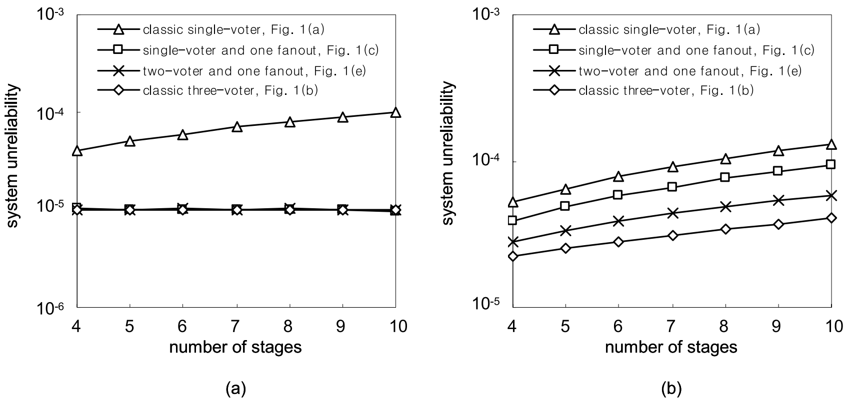

Fig. 3. Effect of the number of stages on reliability: (a) $(1\text{-}R_m)/(1\text{-}R_v) = 1$, (b) $(1\text{-}R_m)/(1\text{-}R_v) = 100$

So far, the number of stages in cascaded TMR modules is assumed to be 100. Fig. 3 compares the four architectures when the length of cascaded TMR modules is much shorter, for two different values of $(1\text{-}R_m)/(1\text{-}R_v)$. The figure shows that the observations we made from Fig. 2 also hold when there are only a few stages in the system. That is, when R_m and R_v are identical, Figs. 1 (c) and 1 (e) are almost as reliable as Fig. 1 (b). When a module is more likely to fail than a voter, Figs. 1 (c), 1 (e), and 1 (b) provide different trade-offs between hardware overhead and system reliability.

4 Conclusion

The classic single-voter and three-voter architectures for cascaded TMR modules have been used for about fifty years. However, this paper shows that there are more practically useful voting structures which provide efficient trade-offs between hardware overhead and system reliability. That is the key contribution of the paper. Specifically, a single-voter architecture with the voter fanout of one is found to be more reliable than the classic single-voter

DOI: 10.1587/elex.4.657
Received September 17, 2007
Accepted September 28, 2007
Published November 10, 2007

ELectronics EXpress

architecture for a reasonably wide range of component reliability. This architecture is almost as reliable as the classic three-voter architecture when the module reliability and the voter reliability are comparable. The two-voter architecture with the voter fanout of one is also practically meaningful. These results are insensitive to the number of stages in the system. Another contribution of the paper is to develop the notion of recovery distance for component failure which allows us to identify efficient fanout topologies with much less effort.

DOI: 10.1587/elex.4.657
Received September 17, 2007
Accepted September 28, 2007
Published November 10, 2007