Chapter 5

Heterogeneous Array Processing

In this chapter, techniques for processing steady-state chemical signals that are the outputs of a heterogeneous array are described. In a heterogeneous array, these signals are the outputs of individual chemical sensors themselves, while in a hybrid array containing both heterogeneous and homogeneous clusters of sensors, these signals are the outputs of homogeneous clusters of sensors preprocessed using one of the techniques described in Chapter 4. Adding heterogeneity to an array of chemical sensors is primarily useful for addressing chemical discrimination. Discriminating among chemicals has long been the most difficult task to accomplish when using microelectronic chemical sensors. In order to perform these discrimination tasks efficiently, it is suggested that signal processing of chemical sensor signals fall within the following three constraints:

- The signal processing should normalize the sensory data sufficiently to minimize the systematic errors caused by changes in the sensor array generated by variations in ambient conditions.
- The signal processing should normalize the sensory data in such a way that the output is relatively insensitive to concentration level changes and drift during the lifetime of these sensors.
- During data normalization, sufficient data should be retained to discriminate among chemicals designated by a particular application.

To meet these three constraints, the signal processing described in this chapter reduces chemical sensor data to a binary representation that expresses the order or ranking of points in a chemical signature between maximum and minimum response. The binary representations generated from experiments here are fairly reproducible across changes in ambient humidity and ambient temperature. Since all of the techniques presented here adapt to global changes in sensory output, the resulting binary outputs are reproducible over changes in concentration and drift. In meeting the third constraint, these binary representations still contain sufficient information to perform chemical discrimination tasks that are typically addressed in the literature and also have practical application.

The heterogeneity of the arrays tested in this chapter is gained by varying operating temperature across the array. Temperature arrays are differentiated from arrays characterized by differences in the actual physical properties of the sensors; the term pseudo-heterogeneous is used to describe these temperature arrays since their structure and range may be altered after fabrication, improving their flexibility and diversity for a variety of sensing applications. Other parameters, such as dopant level and catalyst type, when varied in an array, produce truly heterogeneous arrays, since these parameters, once fabricated, cannot be altered during sensor operation. Whether the array is heterogeneous or pseudo-heterogeneous, it can be designed to be continuous or discrete. Since individual, commercially available sensors are used to test the circuits presented in this chapter, these arrays are discrete in operating temperature. In an integrated system, similar arrays of chemical sensors could easily be made continuous by applying heat to one end of a continuous thin film of chemically sensitive material and monitoring various points along the temperature gradient generated along the film. In this way, any number of discrete points could be extracted as sensory output from a single, gradated heterogeneous thin film.

The signal processing techniques outlined in this chapter can be applied to a variety of pseudo-heterogeneous and heterogeneous arrays of both thin-film and ChemFET-based sensors just by properly choosing the range and type of array parameters. The usefulness of temperature as a pseudo-heterogeneous array parameter is first described below. Circuits that normalize temperature arrays are then described, characterized, and applied to actual arrays of chemical sensors.

5.1 The Temperature Array

When tin-oxide sensors are operated over a range of operating temperatures, they generate a sensory output that varies from one chemical to the next. For example, hydrocarbons such as carbon monoxide tend to react more with tin-oxide at lower temperatures around 150° C than at higher temperatures. Many alcohols, on the other hand, generate a strong response at higher temperatures around 400° C [43]. These variations in response can be used to effectively discriminate among chemicals. Although the exact mechanisms behind the uniqueness of a chemical's signature across temperature are complex and not fully understood in the scientific community, the response across temperature can be approximated by the following equation [55]:

$$S_i = a_1(T_i)x_1^k + a_2(T_i)x_2^k + \ldots + a_n(T_i)x_n^k \tag{5.1}$$

where S_i is the output of a particular sensor or cluster of sensors operating at temperature T_i, the coefficients a_n represent the effects of a particular chemical on a sensor at the temperature T_i, the variables x_n are the concentrations of various chemicals in the sensing environment, and k is a constant dependent on the sensor technology. The coefficients a_n are primarily responsible for creating a unique signature for many chemicals since they reflect the mechanisms behind the reaction of a particular chemical with the sensor surface. An array of sensors operating at N different temperatures would generate N of the these outputs S_i, where the a coefficients are different for every temperature in the array. Fluctuations among sensors in drift and mismatch and saturation effects can affect the response curve, $\boldsymbol{S}$, over temperature of chemical sensor array. The result is that a signature for a particular chemical varies slightly from one response to another (Figure 5.1) [43].

In order to perform robust chemical discrimination, it is desirable to normalize the data in this signature such that discrimination information is retained while variations in sensor performance are not. One way to reduce a chemical signature in this manner involves retaining the order or ranking of individual points in the signature while discarding changes in curvature and steepness along the response curve. While this normalization method does eliminate many of the subtle changes in a chemical signature, it retains sufficient data to discriminate among a wide variety of chemicals. The more closely related two chemicals are, however, the more finely spaced the dimensions of a heterogeneous array of sensors must be in order to discriminate between the two chemicals. When temperature is used as the array parameter, finer spacing can be obtained either by arraying the same number of sensors over a smaller temperature range or by increasing the size of the array. In subsequent sections, various methods of sorting the steady-state outputs of a 10 element, pseudo-heterogeneous array of tin-oxide sensors are presented. The collective nature of each rank-order filtering circuit makes it easily scalable to a larger array when the development of integration technologies for sensors and circuits has matured.

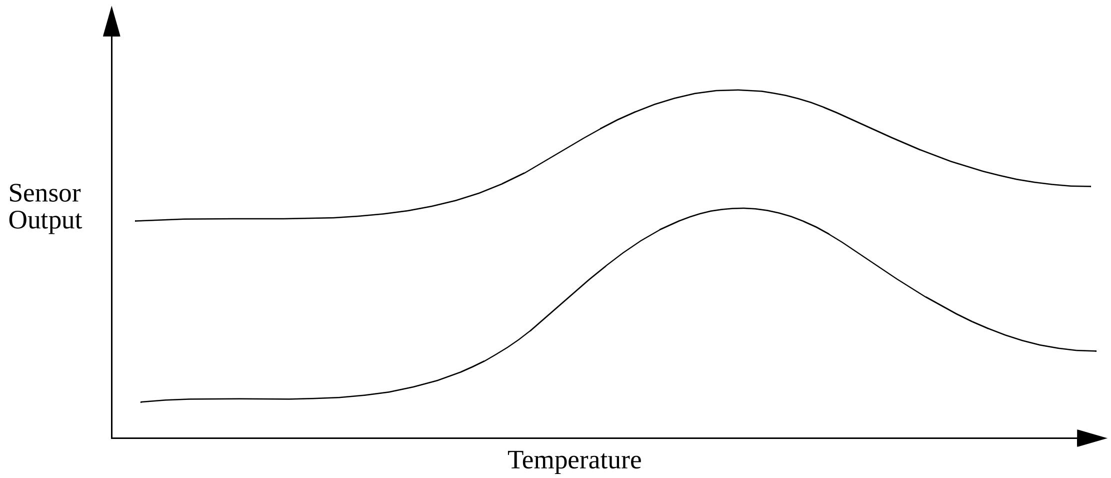

Figure 5.1: Variability in Chemical Signatures

A signature for a chemical, while unique and distinguishable from that of other chemicals, can vary slightly in response to changing concentration and changing ambient condition levels. The above figure shows the tin-oxide response to a typical reducing chemical at two different concentrations (upper and lower curves). Drift over the lifetime of these sensors also causes the response curve for a particular concentration of ethanol to move upward. Eventually, the sensors become no longer useful, as drift eliminates their dynamic range [43].

5.2 Circuit Descriptions

In this section, signal processing circuits are presented that convert a heterogeneous array of chemical sensor output into a compressed binary representation using the following algorithms:

- Adaptive Thresholding: the array is ordered according to a single adaptive point (mean or median value) that shifts with global changes in the array. An array of *N* heterogeneous sensor inputs generates an array of *N* binary outputs that represents the chemical in the sensing environment. Adaptive thresholding provides the most data compression of the techniques presented in this chapter because only a single reference point for ordering is computed; all sensory inputs lying above and below this reference are converted to binary high and low values, respectively.
- Rank-Order Filtering: the array of chemical sensor inputs is arrayed according to rank or place in the array. A single rank-order filtering step might involve the detection of the winner or peak value in the array. Multiple filtering steps detect the second-highest, losing, and similar points in the array; the extent of the ordering is chosen to sufficiently discriminate among the chemicals of interest in a particular chemical sensing application.

Techniques for both adaptive thresholding and full-scale rank-order filtering are described followed by experimental results from the fabrication of these circuits. Circuits are implemented in analog VLSI to threshold an array of sensory input based on an adaptive reference value and to detect single and multiple-ranks in these same arrays. The performance of these rank-order filtering circuits in discriminating among chemicals from a tin-oxide sensor array is then discussed and evaluated in order to evaluate their effectiveness of these circuits for discriminating among chemicals as compared to other methods based on more traditional DSP architectures.

5.2.1 Adaptive Thresholding of an Array of Analog Inputs

In adaptive thresholding, a single reference value is chosen for ordering an array of sensory inputs. Thresholding itself is simply the conversion of a set of analog inputs into a corresponding set of binary outputs, based on the comparison of each input to some reference value. This reference value may be either a fixed or adaptive quantity. Thresholding based on a fixed reference value is accomplished as follows:

$$f(i_n) = \begin{cases} 0 & \text{if} \quad i_n < I_{\text{ref}} \\ 1 & \text{if} \quad i_n > I_{\text{ref}} \end{cases} \tag{5.2}$$

where the quantity I_{ref} is a constant. The output $f(i_n)$ for element n in an array of N elements is a binary low for the input values less than the reference value I_{ref} and high for the inputs greater than I_{ref}. Fixed thresholding using this algorithm is well established in VLSI hardware in a variety of comparator designs and architectures [56]. One of the primary drawbacks of fixed thresholding, however, becomes evident in applications where the background levels in an array of sensory inputs changes while the basic input pattern or image does not itself change. This drawback may be intuitively understood by considering the application of fixed thresholding to image processing applications. Consider the visual scenes shown in Figure 5.2. When evaluating the original image under indoor lighting (Figure 8.1a), the binary output resulting from fixed thresholding is an accurate representation of the primary objects (circle, square) in the image. When the visual scene is moved outdoors (Figure 5.2c), however (Figure 5.2d), the resulting binary output is no longer accurate, since the additional background lighting in an outdoor scene is itself above the reference value used for thresholding.

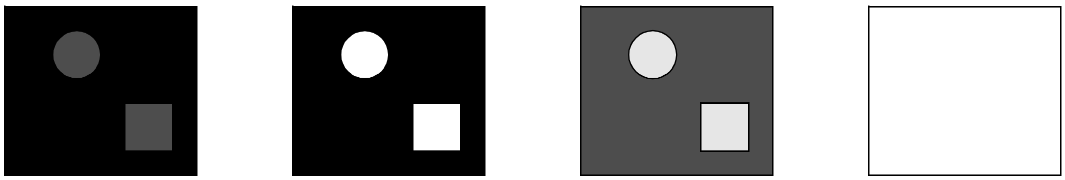

(a) (b) (c) (d)

Figure 5.2: Drawbacks of Fixed Thresholding in Visual and Chemical Image Processing
Fixed thresholding can generate two different outputs for the same basic image. The only difference between the two visual scenes is that (a) the first original image is at a lower background illumination than (c) the other original image. For the constant reference value used in this example, the output of the first image is an accurate one (b), while the output for the second image (d) provides no information about the two objects in the original scene. For a chemical scene or sensing environment, the two original images in (a) and (c) might be representative of two different counteractions of the same chemical across a heterogeneous tin-oxide sensor array.

In terms of chemical sensing, this lack of resilience to global or background offsets may be understood by varying the concentration of a particular chemical in the sensing environment. At a moderate concentration, the binary output generated by fixed thresholding might generate an image similar to the one depicted in Figure 5.2b to represent chemical A. However, if the concentration of chemical A is increased to the point that all sensory output lies above the fixed thresholding value, the resulting binary output might also look like that shown in Figure 5.2d where it is no longer useful for chemical discrimination. Erroneous binary output patterns might also result from changes in such ambient conditions as temperature and humidity and from drift that occurs during the lifetime of each sensor.

In contrast to fixed thresholding, adaptive thresholding is better suited to generate reproducible binary output images because it adapts to global offsets caused by changes in concentration levels, ambient conditions, and similar factors. The adaptive thresholding reference varies with global offsets, generating the same image for a particular chemical regardless of concentration or small variations in ambient conditions. The basic equation that describes adaptive thresholding is similar to that for fixed thresholding:

$$f(i_n) = \begin{cases} 0 & \text{if} \quad i_n < I_{\text{ref}}(i_1, i_2, i_3 \ldots i_N) \\ 1 & \text{if} \quad i_n > I_{\text{ref}}(i_1, i_2, i_3 \ldots i_N) \end{cases} \tag{5.3}$$

The fundamental difference between the fixed thresholding equation (5.2) and the adaptive equation (5.3) above is that for adaptive thresholding, the reference value, I_{ref}, is not a constant; rather, it is a function of the inputs themselves. As background levels (global offsets) in the inputs increase, the reference value used to threshold these inputs also increases. Examples of such variable reference levels are the mean and median values of an array of inputs that have previously been implemented in collective circuits for image processing tasks [57] and are now discussed in the following sections.

Adaptive Thresholding according to the Mean Value in an Array of Inputs

Using the mean value of an array of inputs as the adaptive reference for thresholding, a binary output pattern is generated according to the following relationship:

$$f(i_n) = \begin{cases} 0 & \text{if} \quad i_n < \frac{1}{N} \sum_k i_k \\ 1 & \text{if} \quad i_n > \frac{1}{N} \sum_k i_k \end{cases} \tag{5.4}$$

where outputs that lie above and below the mean value of the inputs are converted to binary high and low outputs respectively. VLSI thresholding circuits that use the mean conversion of (5.4) to generate a binary output pattern have been fabricated, in a standard 2.0µm n-well CMOS process. An implementation of mean thresholding is shown in Figure 5.3.

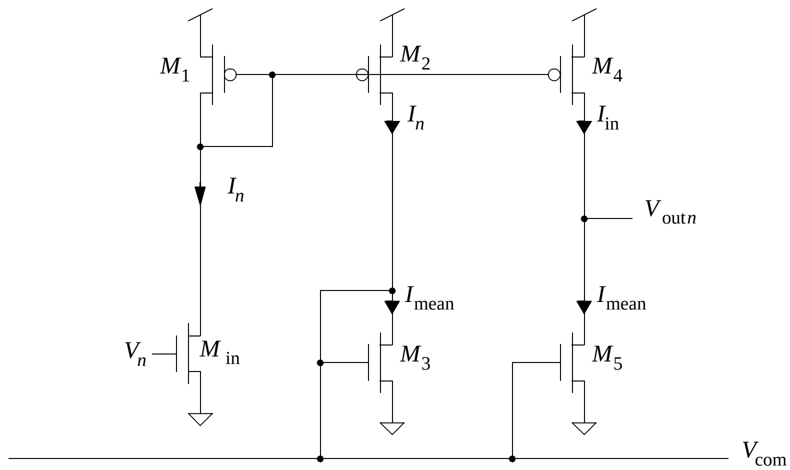

Figure 5.3: The Mean Thresholding Element

Since the common node is attached to the same point in all array elements, the sum of all the input currents I_n must equal NI_{mean} where N is the total number of pixels or array elements. The mean current I_{mean} is then compared with the input current I_n in a single-stage comparator, resulting in a high output (V_{outn}) when the input current is greater than the mean and in a low output (V_{outn})when the input current is less than the mean. The input current I_n can be generated either directly by a chemical sensor or by connecting a sensor output voltage to the gate of the MOSFET M_{in} shown above.

The input current is mirrored to the transistor M_2 through the $M_1 - M_2$ current mirror, where it flows into the global common node V_{com}. The common node is connected to the same point at every element in the array. At V_{com}, Kirchoff's current law must be satisfied as follows:

$$\begin{aligned} \sum_n I_n &= \sum_n I_{mean} = NI_{mean} \\ I_{mean} &= \frac{1}{N}\left(\sum_n I_n\right) = \text{Mean Current} \end{aligned} \tag{5.5}$$

From the relationship described in (5.5), it is evident that the transistor M_3 in every pixel conducts the mean current I_{mean}. The mean current is then compared with the input current I_n at every pixel through the $M_4 - M_5$ gain stage. Since these transistors form a simple, single-stage comparator, the output $V_{\mathrm{out}n}$ is a binary high value if I_n is larger than I_{mean}, and a binary low if I_n is smaller than I_{mean}.

The elegance of implementing mean thresholding in a collective architecture arises from the fact that the mean value is computed in real-time as a function of the interconnect topology; the global common line, V_{com}, provides all the communication necessary among sensing elements to calculate the mean voltage, V_{mean}. No scanning of sensory input to other processing centers off the sensing plane is required; each thresholding element is also small enough that, integrated onto a single substrate with chemical sensors, these circuits will have an insignificant impact on the real-estate required for each sensing element.

Adaptive Thresholding using the Median Value in an Array of Inputs

In a similar manner, adaptive thresholding can also be accomplished by using the median value of the inputs as a reference. If the number of inputs in an array, N, is even, the median of the array is found by arranging the inputs in ascending order and then taking the average of the two middle inputs ($i_{N/2}$ and $i_{N/2+1}$). The inputs are then thresholded according to this median value as follows:

$$f(i_n) = \begin{cases} 0 & \text{if} \quad i_n < \frac{1}{2}(i_{(N/2)} + i_{(N/2+1)}) \\ 1 & \text{if} \quad i_n > \frac{1}{2}(i_{(N/2)} + i_{(N/2+1)}) \end{cases} \tag{5.6}$$

where the output $f(i_n)$ for the nth array element is a binary high and low for inputs i_n that lie above and below the median value, respectively. In the chemical sensor array used for testing in this chapter, the number of inputs is odd (15). To accommodate an odd number of inputs using equation (5.6), the 16th input to the circuitry is tied to ground so that seven of the sensory inputs generate a binary low output and eight sensory inputs generate a binary high output.

A seven-transistor circuit that thresholds an array of inputs according to the median input value in an array of N elements is shown in Figure 5.4.

To understand the operation of this circuit, first assume that the two bias voltages V_{biasn} and V_{biasp} are set such that I_{biasn} is equal to $2I_{\text{biasp}}$ when the bias transistors are both operating in saturation. Kirchoff's current law requires that

$$\sum_{n} I_{\text{biasn}} = \sum_{n} I_{\text{biasp}} \tag{5.7}$$

Assuming that I_{biasn} is twice I_{biasp}, the transistor M_2 must be "turned off" (have a gate voltage equal to 0) in one half of the circuit elements, such that I_{biasn} does not flow in these elements. This restriction ensures that the I_{biasp} sum is equal to the I_{biasn} sum. The selection of which elements turn off I_{biasn} is made via the feedback from the output V_{outn} to M_2. The elements in which the transistor M_2 is turned on correspond to the highest input values in the array. If the input current, I_n, is large, the voltage $V_{\text{out}n}$ goes high, which turns M_2 on, sinking the current I_{biasn} for that element. Similarly, if I_n is small, $V_{\text{out}n}$ goes low, turning M_2 off, thereby preventing that circuit element from sinking I_{biasn}. If more or less than half the pixels attempt to generate binary high outputs, the common node voltage will respond by decreasing or increasing, respectively, until the bias current balance is once again maintained.

5.2.2 Rank-order Filtering of a Sensory Input Array

The concept of thresholding an array of sensory inputs can be extended by further determining the exact location of certain points in the array relative to other points. Thresholding techniques pick only one ordering reference. Rank-order filtering will pick multiple points. The simplest form of rank-order filtering can involve the selection of the maximum (winner-take-all) or minimum (loser-take-all) value in an array of sensors. While these two points are sufficient for performing some sensory discrimination tasks, finer discrimination requires further ordering of points in the sensory array. In these cases, it may be necessary to detect the runner-up values (second-highest and second-lowest) values in an array. Further discrimination capability may be obtained by expanding these basic rank-order filtering techniques to any number of points in a given heterogeneous array. Circuits that perform these rank-order filtering tasks have been fabricated as follows:

- Winner-take-all (peak detection)
- Loser-take-all (minimum detection)

- Multiple rank-order detection implemented in fully-parallel and semi-parallel fashion

Winner-take-all and loser-take-all filtering are the simplest examples of rank-order filtering as they involve detection only of a single rank (maximum or minimum). Runner-up detection extends the rank-order filtering to two ranks and multiple rank-order filtering allows the order of any or all points in an array to be determined.

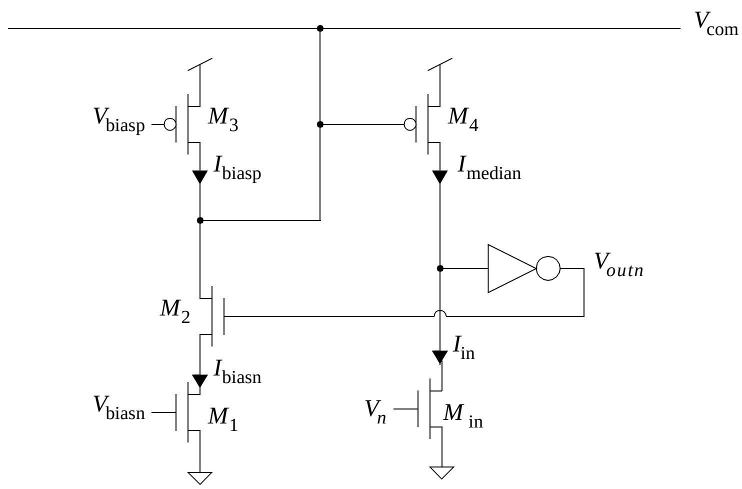

Figure 5.4: The Median Thresholding Element

If the bias currents are set such that $I_{biasn} = 2I_{biasp}$, only half of the circuit elements can sink I_{biasn}, resulting in a high output (V_{outn}) at only half of the elements and a low output (V_{outn}) at the remaining elements The elements that "win" a high output voltage correspond to the highest valued inputs (V_n) in the array. Similarly, the elements that produce a low output voltage correspond to the lowest valued inputs in the array.

Single Rank Detection

The simplest case of rank-order filtering, the winner-take-all, has been used frequently to detect a primary object of interest or target in focal plane processing tasks and detects the peak or maximum value in an array of inputs according to the following relationship:

$$f(i_n) = \begin{cases} 0 & \text{if} \quad i_n < \text{maximum}(i_1, i_2 \ldots i_N) \\ 1 & \text{if} \quad i_n = \text{maximum}(i_1, i_2 \ldots i_N) \end{cases} \tag{5.8}$$

Only the circuit element that corresponds to the peak value in an array of N inputs generates a binary high output. All other outputs remain low. This nonlinear transfer characteristic can be obtained from a winner-take-all circuit (Figure 5.1) implemented in analog VLSI that is very similar to one designed by Lazarro [58].

Each element of the winner-take-all circuit takes an output voltage from a chemical sensor or preprocessed homogeneous cluster of sensors (Chapter 4) as its input voltage. The output of each element is a binary value that is low for all elements in the array except the highest valued input that generates a binary high output. Each node or processing element in the winner-take-all circuit consist of three transistors that are configured to give a highly nonlinear transfer function. The transistors M_2 are all connected at their gates by a single common wire that operates at a voltage V_{com}. This single communication line ensures that all M_2 transistors operate at the same gate-source voltage and that the value of the voltage V_{com} is set to source the largest current in the array generated by the input transistor M_1. All gate voltages at the remaining input nodes V_{n} are smaller in magnitude than this winning value. Thus, their respective M_2 transistors must handle the lower currents generated by the input transistors M_1 by forcing M_2 out of saturation. Assuming a very small drain resistance, the drain voltage for the M_2 transistors in these losing cells will be high for the cells corresponding to the lower input voltages, and much lower for the highest input voltage. The highly nonlinear feedback in the winner-take-all circuit reinforces this mode of operation by allowing only the cell with the highest input voltage to sink the bias current generated by V_{bias}. In other words, the transistor M_3 is turned on for the winning cell and turned off for all remaining cells. The inverters convert these high and low signals to binary values so that the winning sensory input generates a binary high value, $V_{\text{out}n}$, and the remaining inputs generate a binary low value at $V_{\text{out}n}$.

Similar to the winner-take-all circuit, the loser-take-all circuit is a single-stage rank-order filtering circuit that detects the minimum value in an array of inputs and generates an active output at the cell corresponding to this losing input. Loser-take-all filtering is performed in a manner similar to winner-take-all filtering with the exception that the losing or minimum value in an array of inputs, rather than the winner, is now the primary point of interest:

$$f(i_n) = \begin{cases} 0 & \text{if} \quad i_n > \quad \text{minimum}(i_1, i_2 \ldots i_N) \\ 1 & \text{if} \quad i_n = \quad \text{minimum}(i_1, i_2 \ldots i_N) \end{cases} \tag{5.9}$$

.

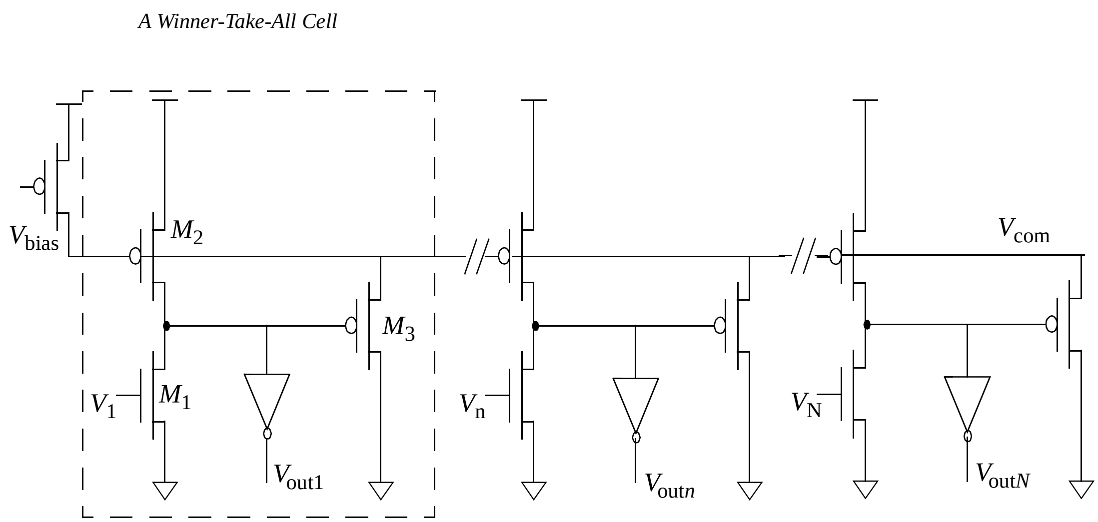

Figure 5.5: The Winner-Take-All Circuit

The simplest case of rank-order filtering is the determination of the maximum value in an array of inputs; such winner-take-all filtering has been applied in various research efforts to focal plane, image processing. The version of the winner-take-all circuit shown above is a slight modification of one designed by Lazarro [58]. V_n and V_{outn} denote the nth input voltage and nth output voltage respectively. V_{bias} is the bias voltage and V_c is the common voltage that is connected to the gate of M_2 at every element in the array. The element that contains the largest input voltage will determine the value of V_{com} and will sink all of the bias current through the transistor M_3. The nonlinear output is achieved when all remaining M_2 transistors are pushed into the ohmic region of operation in order to compensate for their respective smaller input currents. The inverter just converts the winning output to a binary high output signal, V_{outn}, and all remaining outputs to a binary low signal, V_{outn}. The winner-take-all cell forms the fundamental building block for subsequent rank-order filtering circuits described in this chapter.

In loser-take-all filtering, only the output $f(i_n)$, corresponding to the smallest input in the array, generates a binary high output, while all other elements in the array generate a binary low output. This nonlinear transfer characteristic can be generated by using a circuit element that is a simple extension of the winner-take-all circuit. This circuit (Figure 5.6) is a straightforward modification of the winner-take-all circuit of Figure 5.1, where the sensor input voltages, V_n, are applied to the gate of a pFET rather than an nFET. In this way, the highest input or drain current through M_1 is generated by the lowest input voltage V_n.

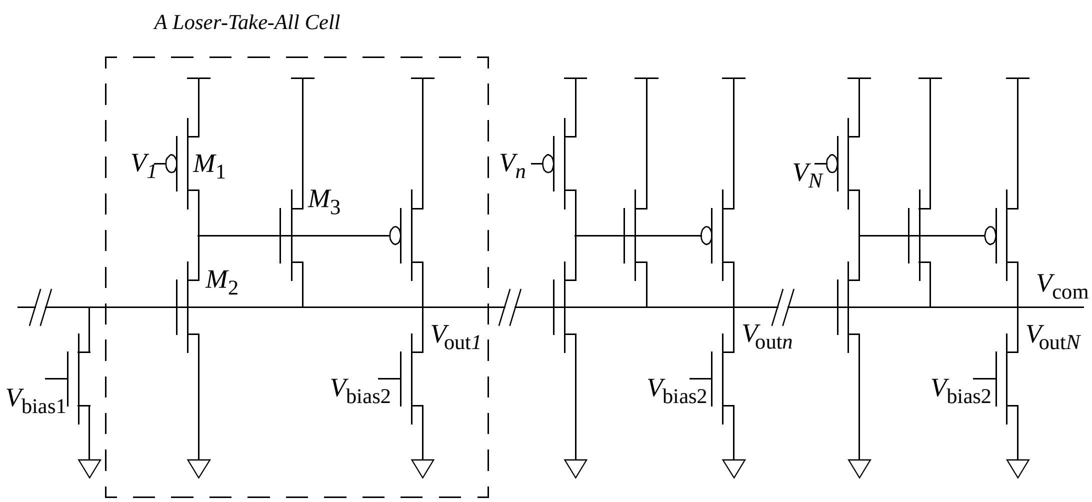

Figure 5.6: The Loser-Take-All Circuit

The loser-take-all circuit is a simple extension of the winner-take-all circuit and detects a single rank (the minimum) in an array of inputs. The lowest or minimum input voltage, rather than highest input voltage, now generates the highest input current and subsequently the active high output voltage.

The winner-take-all cell is the basic building block used for performing multiple rank detection tasks. If a winner is chosen from an array of inputs, the corresponding output from the winning cell can then be used to inhibit another series of winner-take-all elements whose inputs are the same as those for the original winner-take-all competition (Figure 5.7).

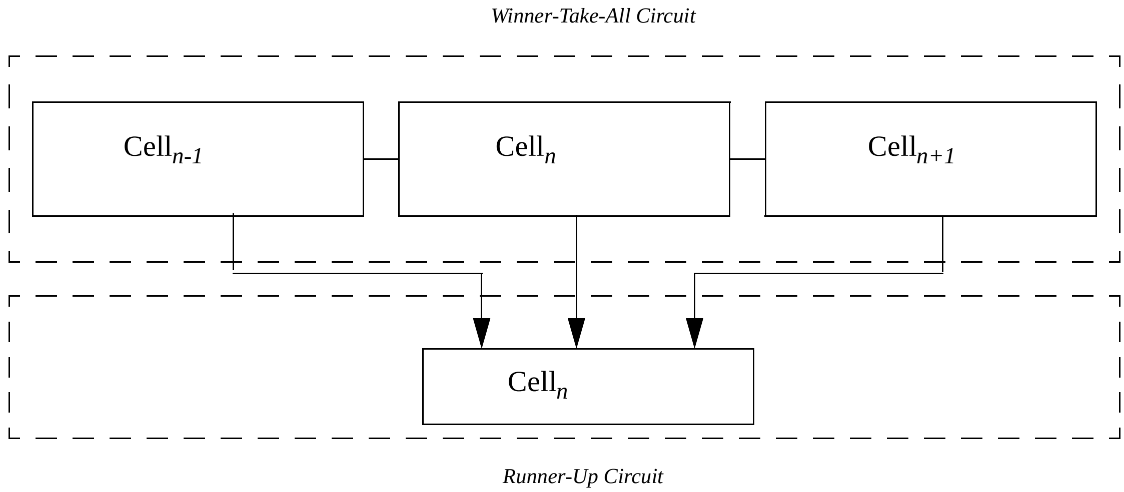

Figure 5.7: Inhibition of the Winner-Take-All Circuits for Multiple-rank-order Detection

In an array of inputs, the winning or maximum input is first detected and an active signal generated at its corresponding output. This output is then used to inhibit the corresponding cell in the runner-up circuit, so that it cannot win the runner-up competition. With the winner inhibited, the runner-up layer selects the second highest input in the array as the winner in this layer. The winning runner-up output then inhibits the corresponding cell in the third-place layer and so on. Inhibition of the *n*th layer may come just from the corresponding output in the (*n*-1)th layer or from the corresponding output and its neighbors in the (*n*-1)th layer as shown above. In this way, the neighborhood around the winner for a particular layer may be inhibited, so that the competition in the following layer selects the true, next highest valued region of interest in the sensing plane.

Once the winner is inhibited in this second series of winner-take-all elements, the active output of the second stage of rank-order filtering corresponds to the second highest value in an array of inputs. A fully parallel implementation of this multiple rank-order filtering scheme is shown in Figure 5.8. The actual detection of ranks below the winner occurs in a similar manner to the winner-take-all selection. For example, in the runner-up (second highest rank) cell, the transistor M_{r4} inhibits the highest valued input (winner) in the input array. The voltage, $V_{\mathrm{out}n}$, that controls whether the transistor M_{r4} generates a small or large current, is low (active) for the winning cell and high for remaining cells. In the non-winning cells, M_{r4} is almost turned off and essentially does not affect the operation of the remainder of the cell. In the winning cell, however, a low voltage $V_{\mathrm{out}n}$ from the original winner-take-all circuit turns its corresponding transistor M_{r4} in the runner-up cell on, thereby robbing current from M_{r2} and forcing the output voltage, $(V_{\mathrm{outr}})_n$, to rise in order to maintain the current balance between the $M_{\mathrm{r2}}/M_{\mathrm{r4}}$ PFET pair and the input transistor M_{r1}.

The output of the runner-up detection circuit can be used in combination with the output of the original winner-take-all competition to inhibit the corresponding cells in the third layer of this circuit so that the third highest valued input in the array can be detected and so on.

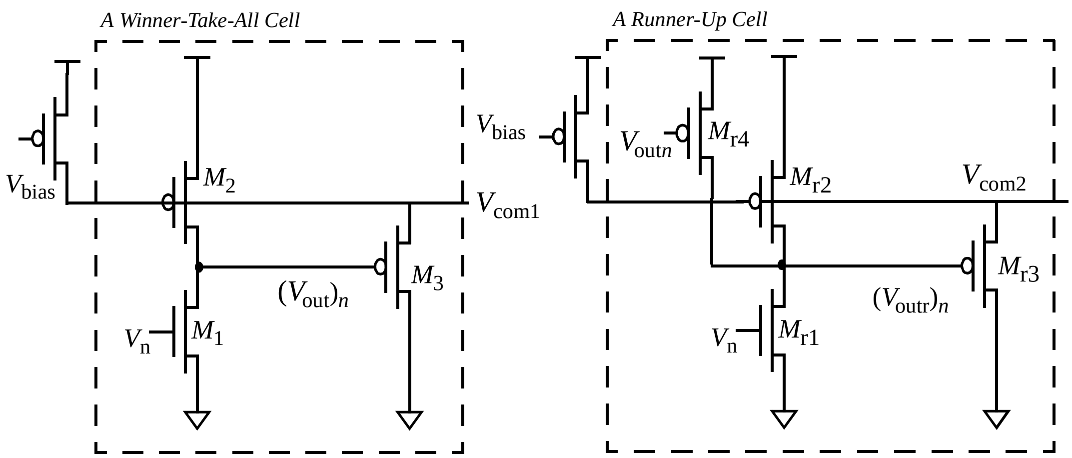

Figure 5.8: Rank-Order Filtering Performed in Parallel

A method for detecting multiple-ranks in an array of analog inputs massively in parallel is shown. Two stages of rank-order filtering are shown here. The winner-take-all cell detects the winner and generates a low output voltage ($(V_{out})_n$) at the winning cell and a high output voltage elsewhere in the array. At non-winning cells, the transistor M_{r4} generates very little current, allowing the runner-up cells to operate in the same manner as the corresponding winner-take-all cell. At the winning cell, however, the transistor M_{r4} is turned on by $(V_{out})_n$, allowing M_{r4} to supply most if not all of the current that M_{r1} requests. This inhibition by M_{r4} prevents the transistor M_{r2} from trying to supply current to M_{r1}. Hence, other cells in the runner-up circuit compete for the new winner selection (runner-up) and control over the common node V_{com2}. V_{com1} and V_{com2} are the common line voltages for the winner-take-all and runner-up cells respectively and the runner-up output voltage, $(V_{out})_{rn}$ is active low. This parallel implementation of rank-order filtering can be extended to the detection of the third highest input in an array by using the outputs of the winner-take-all and runner-up competitions to inhibit the next layer of this circuit.

The inhibition of the winner in an array may also be extended to include the neighbors of the winner, preventing them from winning the runner-up competition as well. In a chemical signature, it is often the neighbors to the winner (e.g. those sensors at nearby temperatures) that generate the second-highest valued sensory output. These neighbors, however, often represent the same general area, or temperature range, of chemical activity. In these cases, neighborhood inhibition in the runner-up cell can be used to detect the true primary and secondary areas of chemical sensitivity in an array of chemical sensors. In actual practice, the inhibition of the winner in the runner-up cell can easily be extended to the sensing elements in the neighborhood of the winner by including an inhibition transistor in parallel with M_{r4} at every element where inhibition is desired. This parallel configuration does decrease the gain of the runner-up circuit by decreasing the equivalent drain resistance of the inhibition portion of the circuit; however, this lost gain may be recovered by increasing the transistor aspect ratios or M_{r4} and M_{r2} or by adding an inverter at the output of the runner-up cells.

Using this parallel rank-order filtering implementation, one series of winner-take-all elements per rank will be needed to detect all necessary ranks for a particular application in the massively parallel manner that is characteristic of these collective architectures. For large arrays of N inputs, however, N inhibited winner-take-all elements will be required to detect all N ranks. If this computation is performed in parallel, the space occupied by all the rank-order filtering elements quickly becomes prohibitively large for sensing plane applications. Since circuit speed is insignificant compared to sensor response speed in these systems, the speed gained by performing rank-order filtering completely in parallel does not justify the extra space required for the fully parallel implementation.

As a result, it is more efficient to implement the complete rank-order filtering computation in a semi-parallel rather than a fully parallel fashion. A semi-parallel implementation of complete rank-order filtering in the sensing plane is shown in Figure 5.9. Rather than processing points in parallel, this rank-order filtering scheme detects N ranks in an array of inputs in (N-1) clock cycles. Before the clock is activated, the winner or peak value in the array is detected in parallel. During the first clock cycle, the runner-up value is detected in parallel; during the second clock cycle, the third highest value in the input array is detected and so on. A single array of N outputs indicates the winner of the competition during each clock cycle. Because of the high gain in this circuit, only one output voltage is active (binary low) during any given clock cycle.

Using only 18 transistors per node in the array, this circuit can be adjusted by varying the clocking scheme in order to detect any number of ordered points in the sensory input. The transistors M_1, M_3, M_4, and M_5 form the same winner-take-all cell that has been discussed previously. The transistors M_6, M_7, and M_8 compare the output current from each winner-take-all cell to a reference current, generating a low output voltage for the winning cell and a high output elsewhere. The selection of the winner is adjusted through successive inhibition of winning cells by the transistor M_2. V_{reset} ensures that after a reset occurs, no cells in the array are inhibited and all transistors M_2 in the array are turned off. At this time, the voltage V_{out} is low for the winning cell and high for all remaining cells in the array. When the clock goes high, M_{10} turns on and allows the value of V_{outn} to pass to $V_{inhibit}$. During the first high cycle of the clock after a reset, the voltage $V_{inhibit}$ goes low for the winning cell and high for all remaining cells. When $V_{inhibit}$ goes low, however, it turns the transistor M_9 off, thereby disabling further feedback from the output voltage V_{outn} to the inhibiting transistor M_2. During the low portion of this clock cycle, the voltage $V_{inhibit}$ is transferred to the gate of M_2, thereby inhibiting the winner or maximum value in the input array.

During the first clock cycle, then, the winner becomes inhibited by a low voltage at the gate of M_2 and can no longer win the winner-take-all competition, resulting in a high voltage at V_{outn}. At this point, the runner-up or second highest value in the array is free to win the competition, resulting in a low output at V_{out} for the corresponding array element. During the next clock cycle, the runner-up is inhibited as the low voltage V_{outn} for the runner-up cell is transferred to the feedback voltage $V_{inhibit}$. The original winner continues to be inhibited as the transistor M_9 remains off in the winning cell, thereby preventing feedback from V_{outn} at the winning cell to the inhibition process. During the second clock cycle, the third highest value in the array is detected and so on through subsequent clock cycles.

This method of successive inhibition can be stopped at any time by resetting the circuit at M_{11}. The transmission gate in the inhibition feedback loop may be eliminated by adjusting the clock speed and/or clock voltage such that during each cycle of the clock, only one cell of this circuit is inhibited; this adjustment is made so that the high cycle of the clock is shorter than the delay of the winner-take-all detection circuitry (transistors M_1 through M_7). A clock whose speed is slower than the propagation delays in this circuitry would allow multiple inhibition cycles to occur during the same clock cycle. A slower clock speed is useful for skipping ordered points in the array, but otherwise can prevent valuable normalized sensor data from being available at the outputs of the rank-order filtering elements.

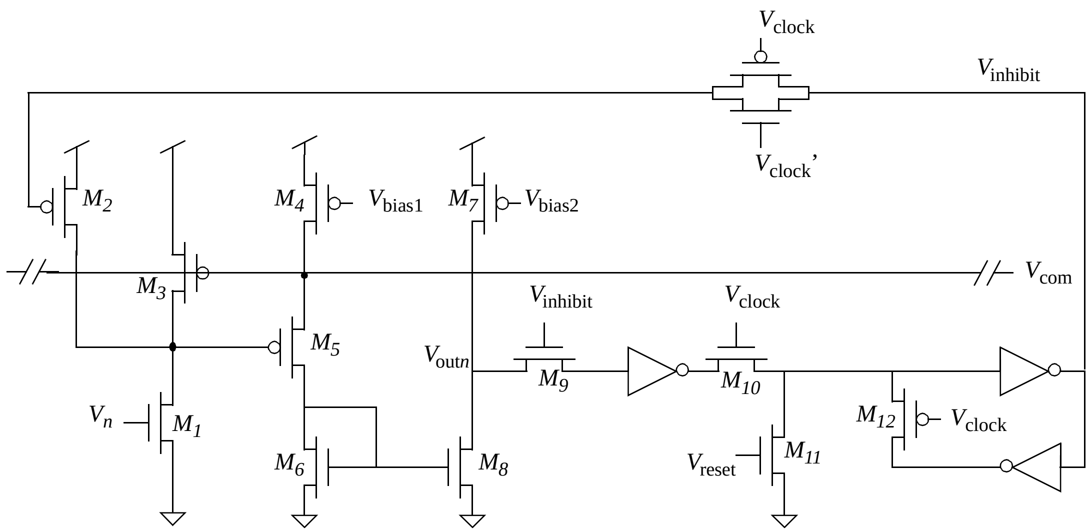

Figure 5.9: Complete Rank-Order Filtering Circuitry

This circuit sorts the points in an array of inputs V_n in descending order. After the circuit is reset, the output V_{out} corresponds to the peak or winning input V_n in the array. During the first high clock pulse, the active low, winning output causes $V_{inhibit}$ to drop low, thereby inhibiting the winning cell through the transistor M_2 during the next low clock cycle, but also shutting off the feedback loop at M_9. With the winning cell inhibited, the circuit will allow the output voltage V_{out} that corresponds to the second highest input value to go low. The process continues until all cells in the circuit are inhibited.

5.3 Circuit Characterization: Experimental Results

Several versions of this family of ordering circuits ranging from the simplest ordering of points found in adaptive thresholding to the complete ordering of points in a sensory array have been fabricated. Circuits have been fabricated in a standard 2.0μm *n*-well CMOS process using the MOSIS fabrication service as follows:

- 16-element array of (mean) adaptive thresholding circuits
- 16-element array of (median) adaptive thresholding circuits
- 4 10-element arrays of rank-order filtering circuits

- 10-element array of winner-take-all circuits
- 10-element array of loser-take-all circuits
- 10-element array of fully parallel rank-order filtering circuits (winner and runner-up only)
- 10-element array of complete, semi-parallel rank-order filtering circuits

Typical experimental behavior, resolution, and error in these circuits are presented in this section. To test the feasibility of these circuits in analyzing actual chemical sensory data, system testing data are also presented in the following section. Overall, these circuits show the potential to perform, with sufficient accuracy for many sensory applications, rank-order filtering tasks on the sensing plane without use less than 20% of the expected sensor area of 4000 μm^2 (see Chapter 7 for details).

5.3.1 Sources of Error in the Rank-Order Filtering Circuits

Two components of error in these circuits are relevant to their performance in chemical sensing systems:

- Offset: the difference between the actual and ideal reference value used for thresholding
- Resolution: the input voltage range needed for an output to switch from fully off to fully on

Especially when it varies across a range of operation for a particular system, offset can cause output patterns to be generated that do not truly threshold, based on the mean and median input values. Circuit offset, or the difference between actual and ideal switching points for the rank-order filtering events, is relevant to the performance of a sensing system because it limits the discrimination capability of the overall system. The range of offset present in these circuits will limit the accuracy of the binary output images. If two points on a chemical signature differ by less than the offset of the signal processing circuitry, they will not be consistently distinguishable. Resolution, on the other hand, is the range of input voltage required for an circuit input to switch from a binary inactive to a binary active value and should logically be less than the resolution of the heterogeneous sensor array from which the circuit inputs are extracted. In chemical sensing systems, resolution impacts the spacing limit for array dimensions and the dynamic range of the system. Both resolution and offset have been quantified in the rank-order filtering circuits and are discussed below.

5.3.2 Adaptive Thresholding: Experimental Results

Arrays of 16 elements designed to threshold an array of analog inputs based on their mean and median value have been fabricated and tested. Results from each type of circuit are presented. Adaptive thresholding is demonstrated as an effective way of choosing a single reference point for ordering an array of inputs while remaining resilient to variations in global offsets caused by changes in concentration levels and in ambient conditions.

The resilience of the adaptive thresholding circuits to changes in global offsets is shown in Figure 5.10. The background levels of the two arrays of inputs are obviously different, yet both the mean and median thresholded outputs do not reflect this change in background level. Both thresholding techniques, becaue of the presence of a global interconnect line, retain the fundamental input pattern only, making them ideal for sensory applications where background level experience a great deal of change during typical operation.

The outputs generated by the mean and median thresholding circuit are fundamentally different. The mean thresholding technique emphasizes (converts to active output) only those inputs that are significantly different from the remaining inputs in the array. However, the median thresholding elements convert exactly half of the inputs in the array to an active output. This difference in operation can be seen also in the responses of Figure 5.10. The mean thresholding elements only convert four of the inputs to an active output, since their magnitudes are much higher than other inputs. The median thresholding elements, as expected, convert eight of the 16 inputs to an active output.

Ideally, in the adaptive thresholding circuits, as an input moves past the reference value (mean or median) in the array, its corresponding input should move from an inactive (binary low) state to an active (binary high) state with infinite gain. In an actual array of these circuit elements, however, finite gain and fabrication variations cause deviations between this ideal behavior and actual behavior. Typical behavior during switching are shown compared with the ideal switching characteristics of the mean and median thresholding circuits in Figure 5.11.

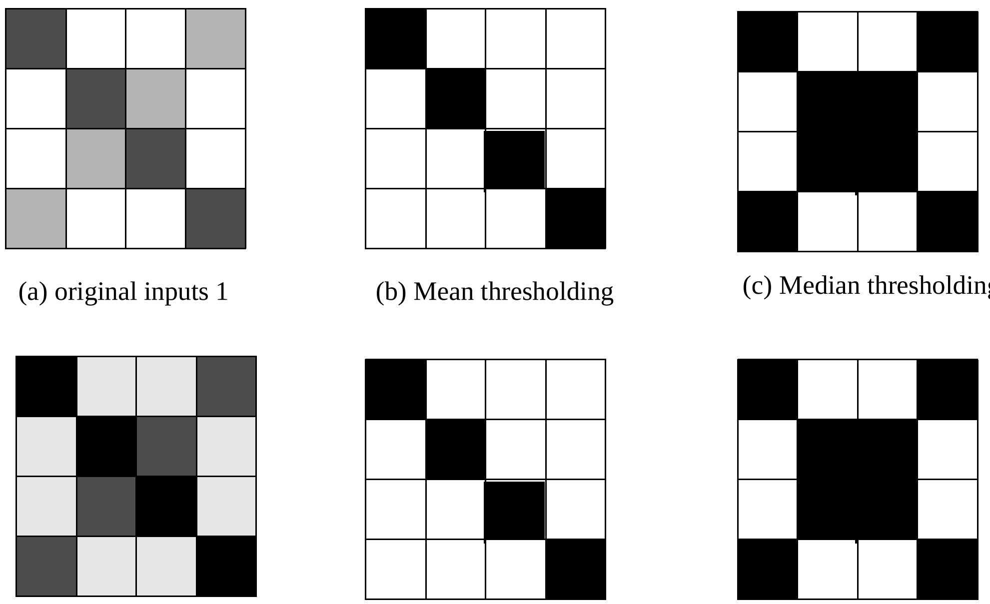

(a) original inputs 1 (b) Mean thresholding (c) Median thresholding

(d) original inputs 2 (e) Mean thresholding (f) Median thresholding

Figure 5.10: Resilience of the Adaptive Thresholding Circuits to Background Changes

The responses of the mean (b, e) and median (c, f) thresholding circuits to two different analog input patterns (a, d) are shown. The only difference between the two input patterns is that the first input pattern has a lower background level than the second pattern. Note that both thresholding circuits retain the fundamental input pattern only, remaining impervious to changes in background levels. The darkest areas in the original input patterns represent the highest input levels and the dark blocks on the output represent active (binary high logic levels).

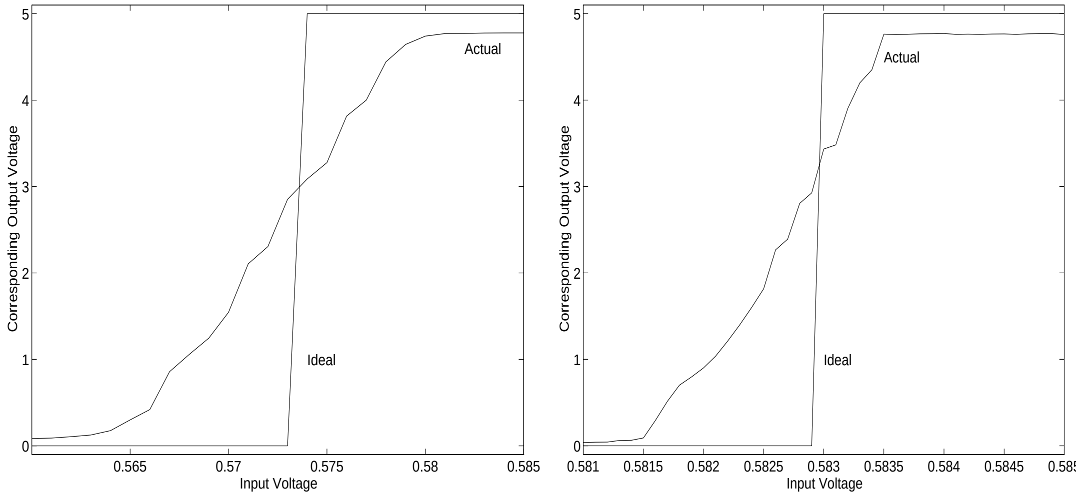

(a) Mean Thresholding (b) Median Thresholding

Figure 5.11: Typical Transfer Characteristics of the Mean and Median Thresholding Circuits
An input (x-axis) in each of the above figures is swept past the (a) mean and (b) median value in an array of inputs and the corresponding output voltage monitored. As the input is swept past the respective thresholding reference of each circuit, the output switches from inactive (low) to active (high). Ideally, this switching would occur exactly at the mean and median value with infinite gain for the mean and median thresholding circuits respectively. Because of transistor mismatch and finite gain in the circuit, however, this ideal transfer characteristic is not achieved and typically limits the accuracy and resolution of the systems in which these circuits are implemented.

When applied to chemical sensing systems, the resolution and offset inherent in the mean and median thresholding circuits is not considered significant since, in subthreshold operation, these sources of error (Figure 5.12) compare with those of a typical microelectronic chemical sensor (tin-oxide) as follows:

- Resolution: 0.02V (mean) and 0.003V(median) compared with 0.1V sensor resolution. The effect of the circuit on the overall resolution of the system is small (<2%).
- Offset: 1.5% (mean) and 0.5% (median) compared with 10% sensor offset [51].

The offset of these circuits is measured as the difference between the ideal switching point (0.5725V in Figure 5.11) and the input voltage that corresponds to an actual output voltage of 2.5 V (approximately 0.52V in Figure 5.11). The resolution, or the input voltage range needed for the output voltage to go from 0 to 5V, is inversely proportional to the gain of each circuit and tends to be better in the median thresholding circuit because of the inverter in the feedback loop that increases the gain during switching. Offset in both circuits tends to stay relatively constant in subthreshold operation (1.5% for mean thresholding and 0.5% for median thresholding) and then to increase as the transistors begin to switch to above threshold operation.

As these chemical sensors are integrated onto a single substrate, resolution and offset in individual sensors will decrease. At this point, it may be necessary to improve the resolution and offset in the mean and median thresholding circuits by employing one or more of the following techniques:

- Improvement in Resolution (gain): increase transistor area or add inverter to output
- Improvement of Offset (mismatch and process variation effects):
  - use of smart layout techniques
  - increase in individual transistor area

Both the mean and median thresholding circuits are small compared with the size of a typical on-chip chemical sensor (less than 10% of total sensor area), so any size increases generated by the above improvements should not have significant real-estate penalties to the overall integrated system.

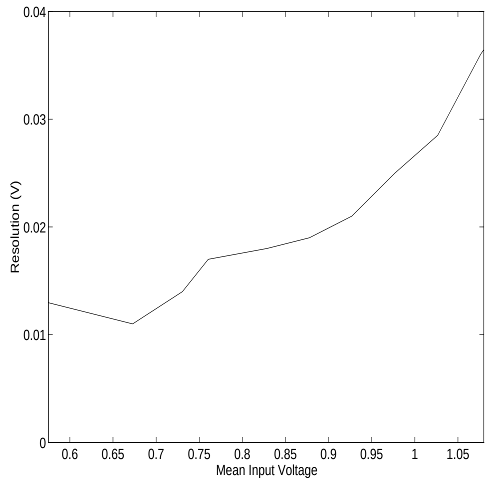

(a) Mean Thresholding, Resolution

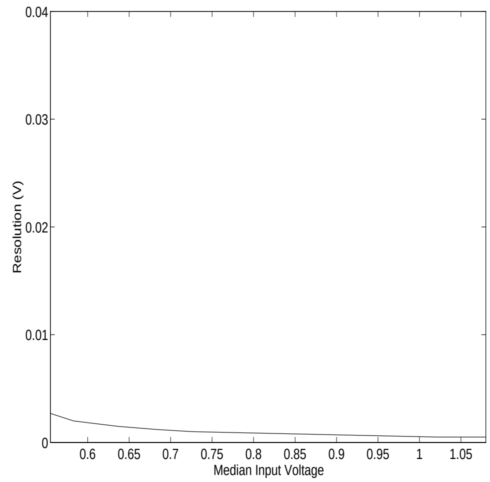

(b) Median Thresholding, Resolution

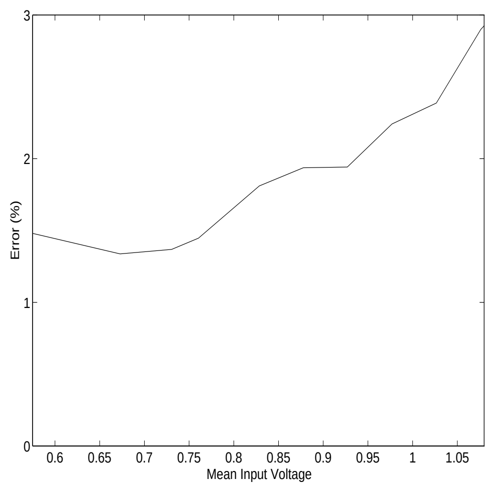

(c) Mean Thresholding, offset

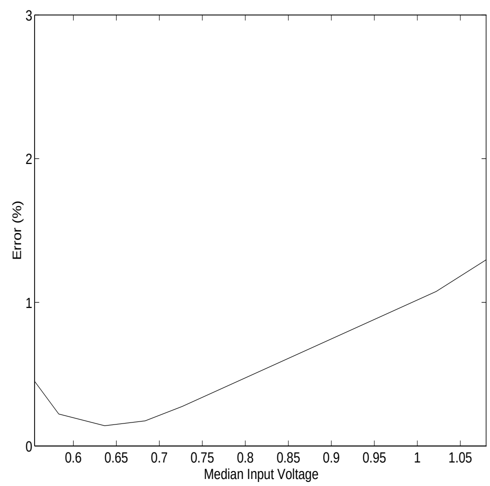

(d) Median Thresholding, offset

Figure 5.12: Offset and Resolution in the Adaptive Thresholding Circuits

The resolution for the (a) mean and (b) median thresholding circuits is shown as well as (c, d) offset in these circuits. Resolution is a measure of the range these two circuits require to switch from fully on to fully off and is inversely proportional to the gain of these circuits. Because of the inverter in the feedback loop, the median thresholding circuit has higher gain during switching and subsequently better resolution than the mean thresholding circuit. Offset, or the difference between actual and ideal switching points in both circuits, tends to increase with increasing mean and median input voltage for the subthreshold operating range shown above.

5.3.3 Rank-Order Filtering: Experimental Results

Arrays of ten elements that determine various degrees of order or rank in an array of inputs are presented in this section. Winner-take-all and loser-take-all filtering are the simplest of these circuits since they locate only the maximum and minimum values, respectively, in an array of inputs. Runner-up filtering extends the winner-take-all to one more level by detecting the peak and second highest values in an array; complete rank-order filtering circuits can, if desired, order every input in the array. Experimental results from various degrees and techniques for rank-order filtering are presented in this section. The application of these rank-order filtering circuits to chemical sensing systems is demonstrated in the next section as these circuits have proven sufficiently robust to produce an reproducible output patterns for certain chemicals.

The winner-take all circuit simply chooses the highest value or peak input in an array and generates a binary active (low) signal at the output corresponding to that element and a binary inactive (high) signal elsewhere in the array (Figure 5.13).

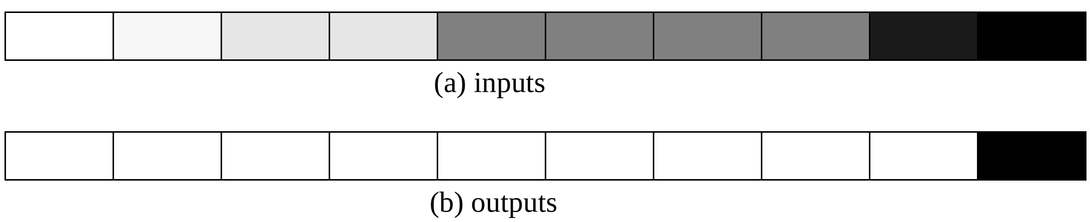

Figure 5.13: Typical Response of the Winner-Take-All Filtering Circuit
The maximum input in an array of inputs wins the winner-take-all competition resulting in a binary active (low) output of the winning element and a binary inactive (high) output elsewhere in the array. The darker blocks in the (a) input field represent higher inputs and the dark block in the (b) output field is active.

A typical transfer characteristic for the winner-take-all circuit is shown in Figure 5.14 and is compared with the ideal switching characteristic of the circuit. Ideally, as an input moves past the current winning input, its corresponding output should change with infinite gain to a binary high value while the output corresponding to the old winning input reverts to a binary low value. Because of finite gain and process variation, however, this ideal behavior is never achieved and like the adaptive thresholding circuits, the errors or differences between actual and ideal behavior for these circuits can be broken into two components: offset and resolution. Typical offset for the winner-take-all circuit is only on the order of 0.16% (of the peak input) and is considered insignificant compared with the estimated 10% accuracy of an individual sensor response. Likewise, the resolution of this circuit is only on the order of 0.001V, a fraction of the 0.1V resolution of a typical tin-oxide chemical sensor [51].

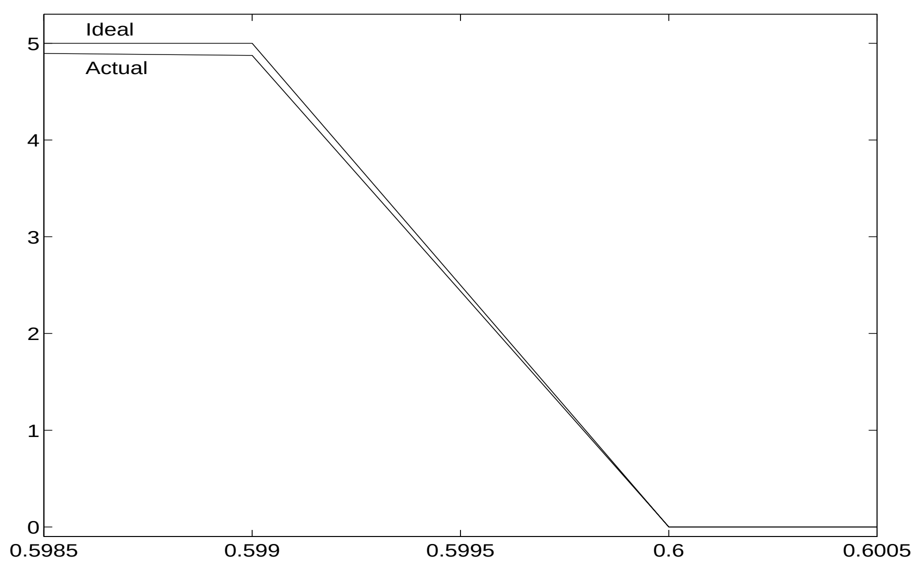

Figure 5.14: Typical Transfer Behavior of the Winner-Take-All Circuit

As an input moves past the current winner (0.6V), its corresponding output moves from high to low while the output of the current winner reverts back to a binary high value. Ideally, this switching point should occur with infinite gain at the exact value of the current winner (0.6V). The output of this circuit is taken before the inverter of Figure 5.1.

The behavior of the loser-take-all filtering circuit is very similar to that of the winner-take-all, with the exception that the minimum value in an array of inputs now generates an active output while all other elements in the array remain inactive. Resolution and offset in the loser-take-all circuit are similar to that of the winner-take-all because the loser-take-all shares a single rank detection scheme similar to the winner-take-all.

Parallel, Multiple Rank-Order Filtering Circuits

The winner-take-all circuit can be replicated and inhibition added to detect any number of ranks in parallel with the winner-take-all competition. For a task requiring M rank-order filtering tasks, these circuit elements generate MN outputs (N is the number of inputs in the array) where M of these outputs are active. A sample output for two stages of the parallel, rank-order filtering scheme is shown in Figure 5.15

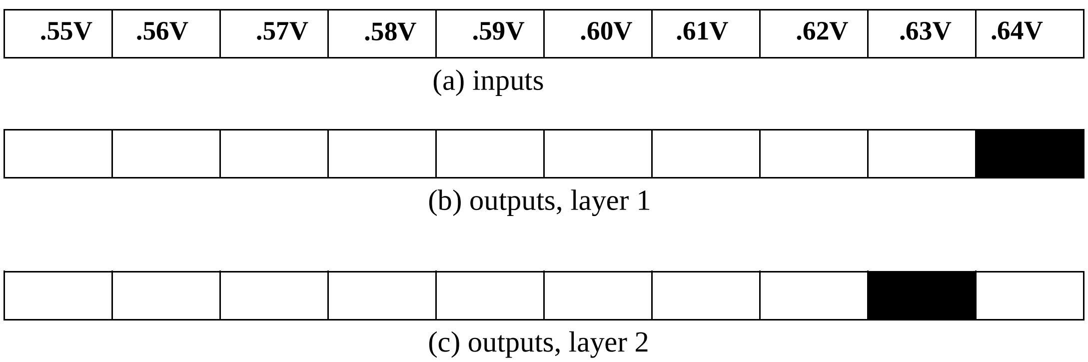

Figure 5.15: Typical Response of the Parallel Rank-Order Filtering Circuit
The maximum input wins the competition in (b) layer 1 of this circuit, generating a binary active output. In the (c) second layer of this circuit, however, the winner is inhibited and the runner-up or second highest input value in the array generates a correspondingly high output. This rank-order filtering scheme can be extended to up to N ranks in an array of N inputs. Parallel rank-order filtering generates a maximum of N^2 outputs, where N of those outputs are active and represent the N ranks in the input array.

The switching behavior of the fully parallel, rank-order filtering circuit (Figure 5.16b) is similar to that for the winner-take-all layer of this same circuit (Figure 5.16a). However, the offset (3.3%) and the resolution (.025V) of the runner-up layer are substantially worse than that for the winner layer of the same circuit. This degradation in performance can be attributed directly to the loss in gain produced by the parallel equivalent drain resistances of the transistors M_{r2} and M_{r4}. (Figure 5.8) in the runner-up layer. The runner-up layer also does not generate a fully binary low value at the output of the second highest input in the array. In order to restore this logic value to a fully off state, an inverter can be added to the output of the runner-up circuit. This addition will also improve the gain, resolution, and offset in the circuit when it is used as part of a larger system. A significant drawback of the fully parallel, rank-order filtering scheme is that the gain of each subsequent layer in the ranking process continues to degrade because of the presence of an additional parallel inhibition transistor at each successive stage of ranking.

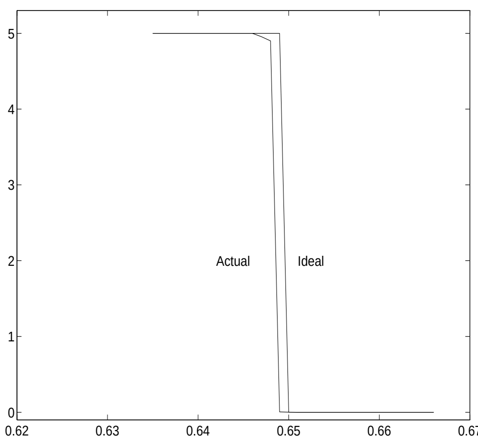

(a) Winner Layer

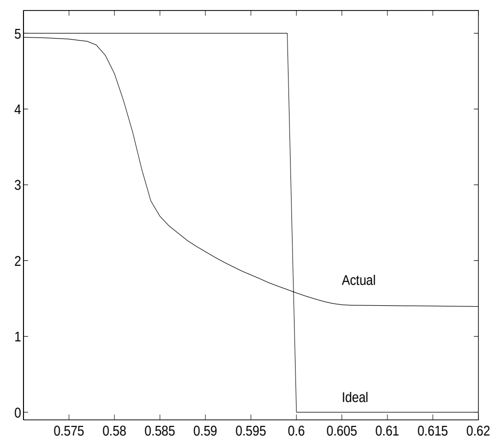

(b) Runner-Up Layer

Figure 5.16: Typical Transfer Behavior of the Parallel Rank-Order Filtering Circuit

Transfer behavior of two stages of the parallel rank-order filtering scheme are shown above. (a) As an input moves past the current winner (0.65V), its corresponding output moves from high to low while the output of the current winner reverts back to a binary high value; ideally this switching point occurs at 0.65V with infinite gain. (b) As an input moves past the current runner-up (0.6V), its corresponding output switches from high to low while the output of the current runner-up reverts back to a binary high value. In the ideal case, this switching point occurs at 0.6V with infinite gain.

Alternatively, rank-order filtering of an entire array of inputs can be performed in a semi-parallel rather than fully parallel fashion. A ten element array of these rank-order filtering elements (Figure 5.9) has been fabricated in a 2.0 μm, *n*-well process (using MOSIS) and has sufficiently robust resolution and minimal offset for application to chemical sensing applications.

As described in Section 5.2.2, the rank-order filtering circuits can also order points in an array of inputs in a semi-parallel fashion. Prior to the first clock cycle, the circuit acts simply as a winner-take-all circuit, generating a binary active output only at the output that corresponds to the peak or maximum input in the array. During the first clock cycle, the circuit acts as a runner-up circuit, generating a binary active value only at the output that corresponds to the second highest or runner-up value in the array. This process continues during the third clock cycle as the third highest input value is selected and so on through all points in the array. Since the selection of the ordered point is controlled by the clock cycle, supplemental signal processing may select any number and rank of points to observe and save for subsequent sensory discrimination tasks.

A typical output for the rank-order-filtering circuit is shown in Figure 5.17 for the first four highest valued inputs in an array of ten inputs.

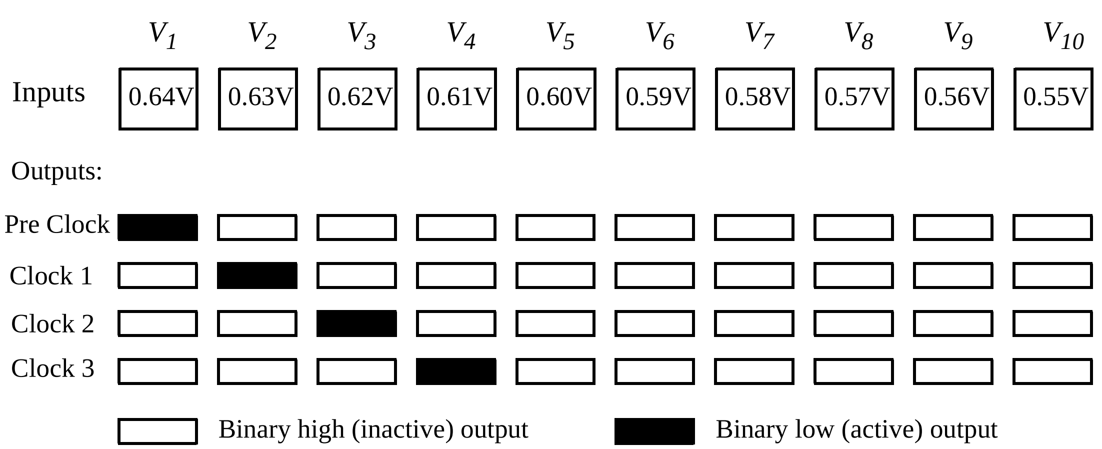

Figure 5.17: Typical Output of the Semi-parallel Rank-order Filtering Circuit

The inputs in the above array are set in descending order from input 1 through 10. The rank-order filtering circuit activates output 1 before the clock becomes active, output 2 during the first clock cycle, output 3 during the second clock cycle, and output 3 during the third clock cycle. If only certain ranks are desired in this filtering process, the read-out circuitry can evaluate output only during the desired clock cycles.

During each clock cycle, only a single output is active, indicating the winner of the current ranking competition. Performing rank-order filtering in this semi-parallel fashion prevents the need for *N* parallel ranking circuits in an array of *N* inputs, which would quickly become prohibitive in cost and real-estate when implemented collectively on the sensing plane with chemical sensors.

As in other rank-order filtering circuits discussed in the previous section, the characteristics of these circuits that are relevant to chemical sensing systems are the offset and resolution associated with each ranking event. Typical resolution and offset for the first four ranks of the semi-parallel rank-order filtering circuit are shown in Figure 5.18. In subthreshold operation, experimental resolution values for resolution and offset are comparable across an operating range of 0.50 to 0.85V, typical for an array of sensors operating at a 1V power supply.

Typical values for resolution and offset are listed in Table 5.1. Resolutions on the order of millivolts and offset in the tenths of percent minimize the impact of the errors in the semi-parallel, rank-order filtering elements on the performance of a typical chemical sensing system. Thus, these circuits are considered robust for processing sensor outputs for chemical discrimination. In the following section, the effectiveness of both rank-order filtering and thresholding elements is evaluated for chemical discrimination in an array of tin-oxide sensors operating at different temperatures.

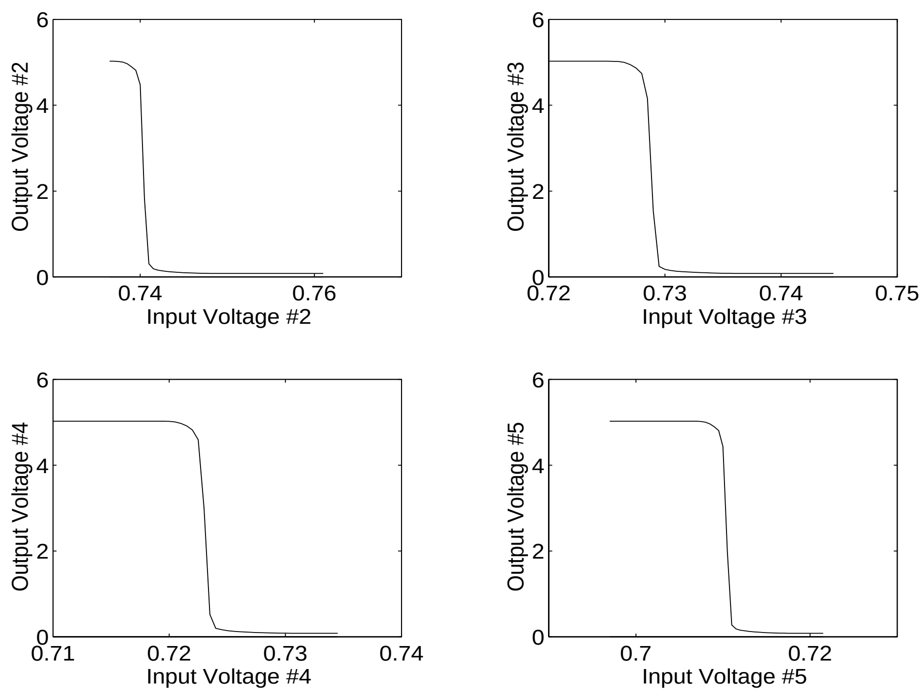

Figure 5.18: Typical Behavior of the Semi-Parallel Rank-Order Filtering Circuit

In the above figures, the input voltage is the winning input in the array at 0.74V. Input 2 is swept past this winning input until its corresponding output switches from high to low, indicating it has won the winner competition. Input 2 is then returned to 0.73V, the runner-up value, and Input 3 is swept past it until its corresponding output voltage (Output #3) wins the runner-up competition. Figures (c) and (d) represent similar switching behavior for the third highest (.72V) and fourth highest (.71V) valued inputs in the array.

TABLE 5.1: Resolution and offset in the Semi-Parallel Rank-Order Filtering Elements

Filtered Rank	Resolution (V)	Offset (%)	Sensor Resolution	Sensor Offset (%)
Maximum value	0.0018	0.125	.1V	10
Second highest value	0.002	0.174	.1V	10
Third highest value	0.0022	0.226	.1V	10
Fourth highest value	0.0025	0.262	.1V	10

5.4 System Testing

In this section, experimental results for the application of the order detection circuits to arrays of heterogeneous chemical sensors are presented. In all system testing, the sensor technology is the commercially available, tin-oxide Taguchi sensor (Figaro Eng). When these discrete sensors are replaced with integrated sensors in future research, the discrimination capability of the overall system should only improve, as fabrication mismatch and process variation will decrease when sensors are fabricated on the same rather than different substrates. In order to test the feasibility of the rank-order filtering architectures for chemical discrimination, however, these circuits have been tested them on the following systems:

- Adaptive thresholding: tested on a heterogeneous array of tin-oxide sensors differentiated by operating temperature and catalyst type.
- Complete Rank-order Filtering: tested on a pseudo-heterogeneous array of discrete tin-oxide sensors: differentiated by operating temperature.

5.4.1 Experimental Set-Up

System testing of the rank-order filtering circuits is performed in a similar manner to the testing of the homogeneous arrays described in Chapter 4. The basic experimental set-up for these experiments is shown in Figure 5.19. Each tin-oxide sensor (Figure 5.19a) contains its own on-board heater. The heater temperature is maintained by a buffer operating at a constant voltage which is consistent with the desired temperature according to specifications provided by the manufacturer. The buffer input is controlled by a variable resistance voltage divider. The tin-oxide sensors used for these experiments consist of the following:

- Adaptive Thresholding
  - Types of Sensors: 5 TGS822, 5 TGS813, 5 TGS880 (Figaro Eng)
  - Sensor Sensitivity: TGS822 (reducing alcohols), TGS813 (ammonia), TGS880 (CO)
  - Temperatures of Sensors: 5 temperatures evenly spaced between 320°C and 360°C
  - Heater Control of Sensors: 5 voltages evenly spaced between 4.1 and 5.0V
  - Total Number of Sensors: 15 (fully heterogeneous)
- Multiple Rank-order Filtering
  - Types of Sensors: 9 TGS822
  - Sensor Sensitivity: TGS822 (reducing alcohols)
  - Temperatures of Sensors: 9 temperatures evenly spaced between 125°C and 485°C
  - Heater Control of Sensors: 9 voltages evenly spaced between 1.9 and 6.3V
  - Total Number of Sensors: 9 (pseudo- heterogeneous)

For the fully heterogeneous array used to test the adaptive thresholding circuits, one buffer is used to control three heaters each or one heater for each type of sensor. In this way, variations in the heater input for a particular operating temperature are minimized. Because the actual temperature of each sensor is independently controlled by the on-board heater, each sensor output is sensitive not only to variations in the actual sensor surface but also to variations in the heaters themselves. The outputs of the rank-order filtering circuits must then remain reproducible across changes in the sensor and in the heater properties over the lifetime of the sensors. The outputs of each sensor array are connected to the inputs of the appropriate processing circuits and the circuit outputs monitored by several source-measurement units via an IEEE-488 interface and Unix-based workstation. All of the sensors are allowed to stabilize for a week at the desired operating temperature before testing is performed.

Various reducing chemicals are introduced into a large chamber (Figure 5.19b) and allowed to evaporate into the chamber. A valve between this chamber and a smaller testing chamber is then opened allowing the gas to diffuse into the environment of the sensor arrays. A fan inside the testing chamber keeps the gas well mixed and evenly distributed.

A large amount of the gas is first introduced into the testing chamber, allowing the sensors to approach saturation and the thresholding and rank-order filtering circuits to be tested in saturation mode. Each test begins with a high concentration of a single gas in the testing chamber. After the sensor arrays stabilize in response to a particular chemical, the testing chamber is gradually aired, while the outputs of the heterogeneous processing circuits are continuously monitored for reproducibility.

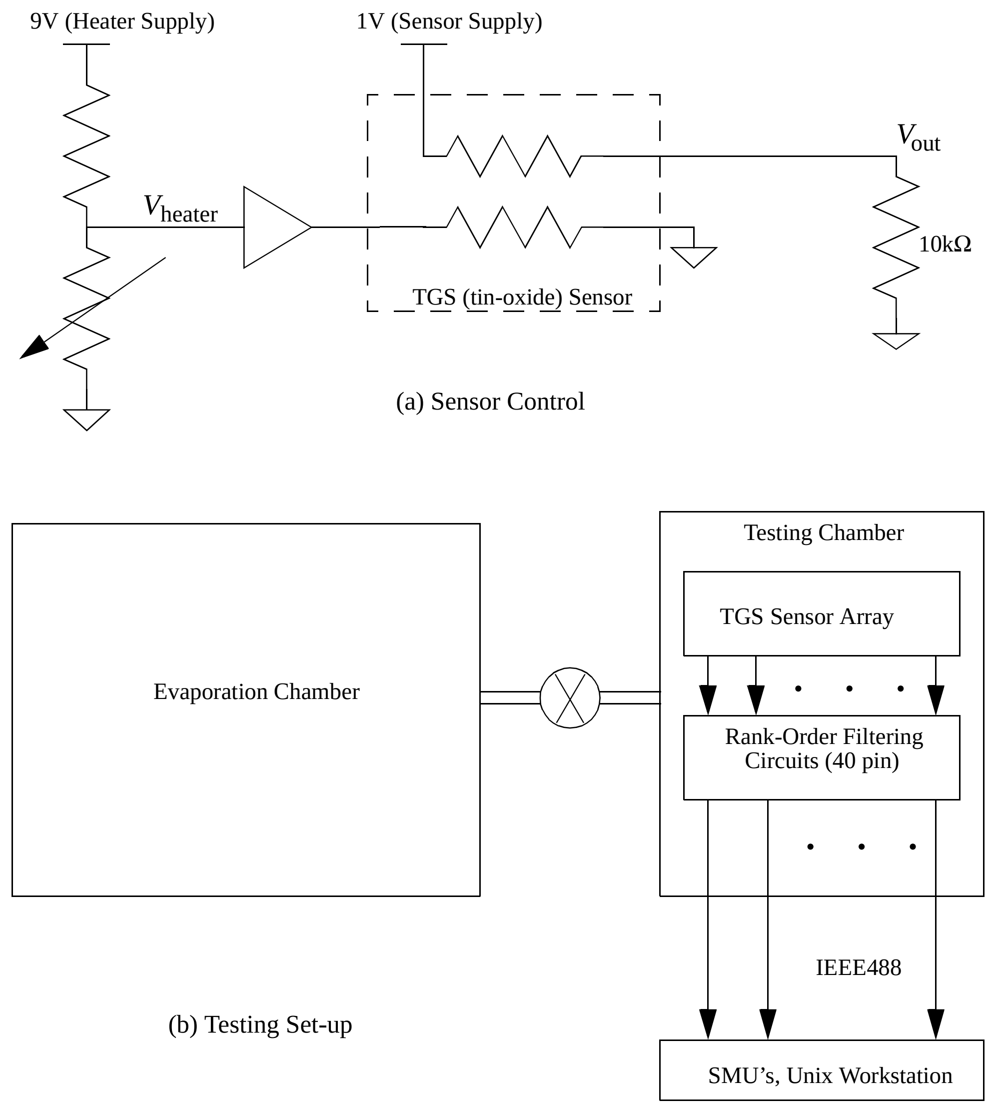

Figure 5.19: Experimental Set-up for Testing the Rank-Order Filtering Circuits

Shown above is the testing set-up for evaluating the performance of the rank-order filtering circuits on an array of tin-oxide sensors. Each sensor (a) consists of a resistive heater and a chemically sensitive resistor. The sensor output voltage is then taken across a 10kΩ load resistor. A chemical is introduced into the (b) evaporation chamber and allowed to evaporate and diffuse throughout the chamber. A valve between the two chambers is opened, allowing the evaporated gas to move into the testing chamber where it is sensed and processed. The outputs are monitored by various source-measurement units and a Unix-based workstation via an IEEE-488 interface.

The sensor outputs are monitored for approximately 30 minutes in both saturation and non-saturation mode. The concentration of each chemical tested decreases during this time period; however, because concentration produces systematic offset in the sensor outputs, the processing circuits should continue to produce the same output patterns. After the 30 minute testing period, the testing and evaporation chambers are aired for at least 30 minutes before the next test is performed. In the following sections, results for both adaptive thresholding and rank-order filtering of the tin-oxide sensor arrays described here are presented.

5.4.2 Adaptive Thresholding System Testing

The mean and median adaptive thresholding elements have been tested at a system level on an array of Taguchi sensors (Figaro TGS822, TGS813, and TGS880). Each of these types of sensors contains a different catalyst diffused into a tin-oxide substrate; although the exact catalyst is not known because of proprietary concerns, these catalysts are known to enhance the sensitivity of the TGS822, TGS813, and TGS880 to alcohols, ammonia, and carbon monoxide, respectively. Each of the three heterogeneous arrays in turn consists of five of each type of sensor, operating at a different temperature, for a total of 15 sensors operating at five temperatures between 320° ×C and 360° ×C.

A heterogeneous array is used in the system testing of the thresholding circuits for two reasons. Heterogeneity broadens the selectivity of the array as a whole sufficiently to discriminate across a wide range of reducing chemicals. Furthermore, the compression of sensory data performed by the thresholding operating is significant and retains minimal information for discrimination. While this compression is useful for minimizing the effect of communication bottlenecks at the chip i/o level, the information retained after the compression for a small array, such as the one used here for testing, has proven sufficient for discriminating among families of chemicals. Finer discrimination capability within the same family of chemicals such as alcohols or closely related families must use either a larger array or a normalization technique that retains more information for discrimination, such as the more complete rank-order filtering also described in this chapter. Despite the large reduction of sensor information inherent in adaptive thresholding, these circuits have nevertheless proven quite useful for discriminating among broad families of chemicals.

A typical response of the chemical sensor array to adaptive thresholding circuitry as described above is shown in Figure 5.20. The output for a particular chemical is a binary chemical image which is often distinguishable from output images generated from the sensory inputs of other chemicals. Each block of the output corresponds to a specific type of sensor (TGS822, TGS813, or TGS880) operating at a specific operating temperature *T.* The gray and white blocks correspond to a binary low output and a binary high output, respectively, from the thresholding hardware.

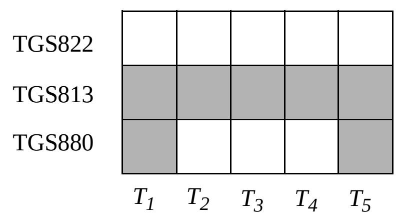

Figure 5.20: Typical Output Response

A typical binary output pattern from the thresholding circuitry in response to sensory input from a heterogeneous array of tin-oxide sensors is shown. The gray and white boxes correspond to sensors that generate a binary high or binary low output at the corresponding circuit outputs respectively.

Mean Thresholding of Heterogeneous Arrays

Thresholding an array of sensor inputs according to the mean value tends to emphasize the outstanding points in the image. In these chemical sensing systems, these outstanding points correspond to the temperatures of maximum reaction between a particular chemical and tin-oxide. These points of maximum sensitivity for a single chemical are known to vary across temperature for different chemicals. Unfortunately, within a mixture of chemicals the response across temperature does not vary linearly. However, some discrimination within a mixture of chemicals can be performed using a heterogeneous array of sensors, since the sensitivity to groups of chemicals can vary significantly from one heterogeneous point in the array to another. As a result, in the thresholding of this heterogeneous array of tin-oxide sensors, it is possible to perform rough sensory discrimination across families of chemicals for single chemical testing and some detection of single chemicals within a mixture of chemicals.

Binary chemical images for six reducing chemicals generated by mean thresholding of this heterogeneous array of tin-oxide sensors is shown in Figure 5.21.

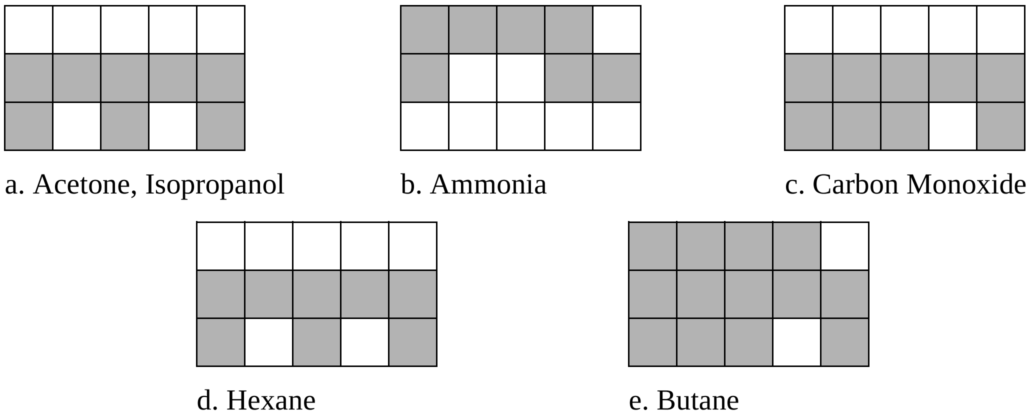

Figure 5.21: Output Patterns for Mean Thresholding of a Chemical Image

A binary output pattern from the mean thresholding process is shown above for six different reducing chemicals. Note that the only patterns that are not distinguishable are those for acetone and isopropanol. Discrimination among these chemicals requires either an increase in the size of the sensor array or finer spacing of temperatures in the array. Note that two of the sensors in the above images always generate a binary high output; these sensors have a baseline value nearly twice that of the other sensors in the array. In practical chemical sensing systems, it is recommended that sensors be matched within 50% of each other to generate fully useful, thresholded output patterns.

As expected from the adaptive nature of the mean thresholding technique, these output images remains constant over a wide range of concentrations (500-5000 ppm). The only chemicals tested here that are not distinguishable are the related chemicals, acetone and isopropanol. Strong chemicals such as acetone and ammonia also generate the same output image as those shown in Figure 5.21 in mixtures containing small amounts of hexane and ethanol. While this array has broad selectivity for discriminating among families of chemicals, more resolution in the heterogeneous array would be required to enhance discrimination capability among more closely related chemicals.

Median Thresholding of Heterogeneous Arrays

Binary chemical images for six reducing chemicals in response to median thresholding of the heterogeneous array of tin-oxide chemical sensors are shown in Figure 5.22. For five of the chemicals tested, the resulting binary image is distinguishable from the remaining chemicals. The images for acetone and isopropanol, however, are identical.

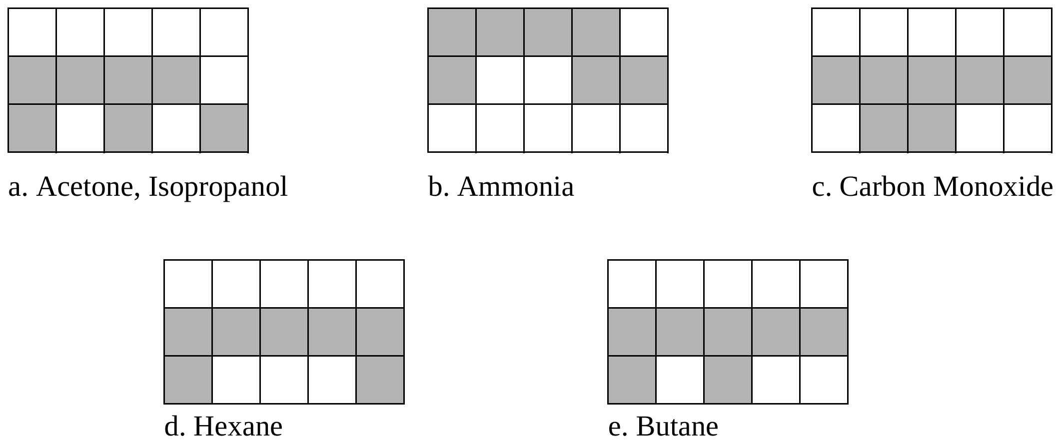

Figure 5.22: Output Patterns for Median Thresholding of a Chemical Image

A binary output pattern from the mean thresholding process is shown above for six different reducing chemicals. Note that the only patterns that are not distinguishable are those for acetone and isopropanol. To discriminate among these chemicals, the size of the array (heterogeneity) could be increased or the temperatures (pseudo-heterogeneity) in the array could be more finely spaced.

Again, the adaptive nature of the median thresholding algorithm suggests that these binary chemical images should remain constant over changes in global parameters across the array such as concentration, drift, or ambient conditions. Indeed, the images shown in Figure 5.22 remain constant over a wide range of concentrations (500-5000 ppm) and over humidity changes and baseline changes caused by drift in the aging sensors. As the concentration of a particular chemical changes, the median shifts also, allowing the output image to remain constant. This adaptability also makes the array output resilient to fluctuations in the sensing environment caused by humidity, ambient temperature, and similar factors.

Large degrees of mismatch, however, do affect the outputs of the median thresholding hardware. In this experiment, sensor #14 (TGS880 at temperature T_4) had a baseline conductivity nearly ten times that of the remaining sensor resistances and as a result, generates a binary high output in response to all chemicals tested. For variations in the baseline conductivity as great as 66%, the median thresholding hardware is not affected and produces unique chemical images in response to different chemicals tested. To maintain the reproducibility of the output images, this hardware should be used in arrays whose baseline conductivity varies no more than 50% from minimum to maximum baseline value in the array.

Median thresholding is fundamentally different from mean thresholding in its application to sensing systems. Since median thresholding of chemical sensor outputs produces a 50/50 high/low image regardless of the quality or quantity of chemicals present, this method of thresholding can provide information about smaller components of a chemical mixture in the sensing environment. For example, in the heterogeneous arrays used for the thresholding experiments, if two very different chemicals such as carbon monoxide and ammonia are present, they will generate a strong response in the TGS822 and TGS880, respectively. In a mixture, one of these two chemicals may have a much stronger concentration than the other; yet, in the appropriately sized array, sensitivity to both chemicals can still be observed by monitoring a thresholded output based on a local median value in the TGS822 and TGS880 sensors.

5.4.3 System Testing: Rank-Order Filtering

In this section, experimental results for rank-order filtering are presented for a nine element pseudo-heterogeneous array where operating temperature varies as a function of a sensor's position in the array. The sensor used in this array is the Figaro TGS822, a tin-oxide film that has been modified to be particularly sensitive to various types of alcohols. The degree of rank-order filtering performed on an array is determined by the discrimination capability required of the sensing system. Winner-take-all, loser-take-all or runner-up filtering can be used in systems where simple discrimination among two or three major chemicals or families of chemicals is desired. More complete rank-order filtering, on the other hand, can be applied to systems where more complex discrimination capability is needed. A sample output for the rank-order filtering system is shown in Figure 5.23. In a fully parallel rank-order filtering circuit, all outputs are available simultaneously. In the semi-parallel rank-order filtering circuits (Figure 5.9), however, only one layer of the output response pattern is available at any given time.

Results for the various degrees of rank-order filtering are presented in Figure 5.24. Winner-take-all filtering is just the first layer in these responses and loser-take-all is the last (bottommost) layer

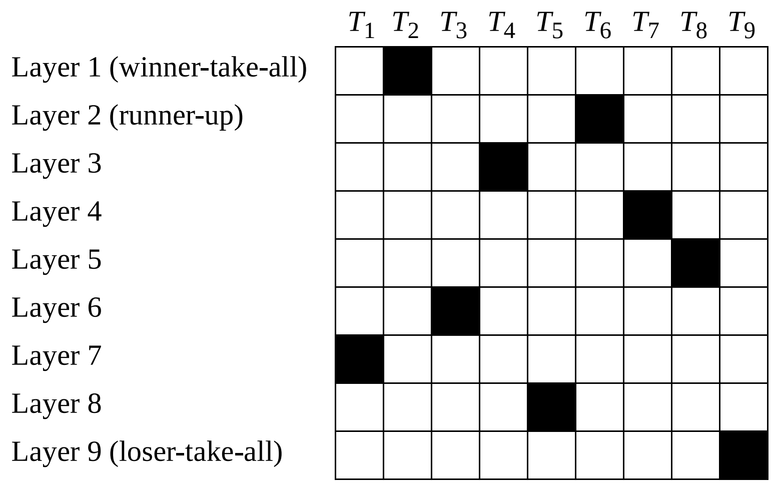

Figure 5.23: A Sample Output of Rank-Order Filtering of a Pseudo-heterogeneous Array
The output pattern above is a sample response to a chemical produced by rank-order filtering of a nine element, pseudo-heterogeneous sensor array operating at temperatures T_1 through T_9. If the winner-take-all filtering circuit is used on this array only layer 1 is generated as the output response. If fully parallel filtering is used on this same array, all output layers are generated simultaneously. If the semi-parallel rank-order filtering circuits are used, Layers 1 through 9 are generated in semi-parallel fashion: layer 1 prior to the first clock cycle, layer 2 during the first clock cycle, and so on. The dark squares in the response pattern above represent the winner of that layer's competition.

of filtering. Part or all of these layers may be used to perform chemical discrimination. In these results, where all nine sensors in the array were ranked, it was possible to discriminate among all of the chemicals tested (ammonia, ethanol, isopropanol, and carbon monoxide). Since ethanol and isopropanol are closely related members of the same family (alcohols), It is certainly feasible to use the rank-order filtering process to perform more subtle discrimination tasks.

5.5 Applications

In this chapter, a series of order detection circuits have been presented for processing heterogeneous arrays of chemical sensors in order to perform chemical discrimination tasks that have not been possible with single chemical sensors. In this section, suitable applications for each of these order detection circuits are evaluated.

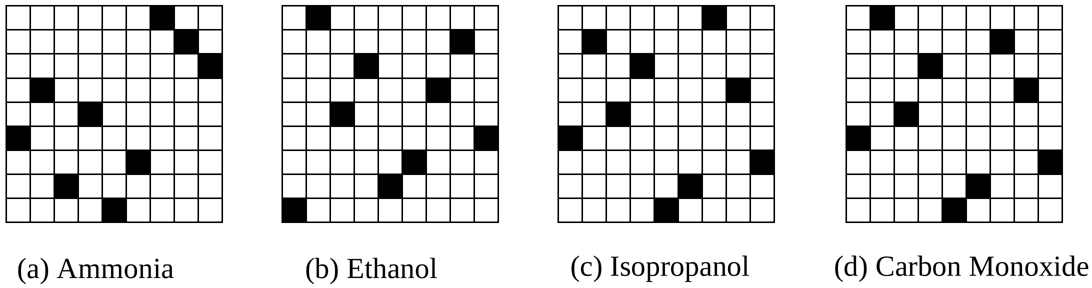

Figure 5.24: Output Patterns for Four Reducing Chemicals Using Rank-Order Filtering
Above are shown the output patterns for four reducing chemicals (a) ammonia, (b), ethanol, (c) isopropanol, and (d) carbon monoxide after rank-order filtering of an array of 9 chemical sensor outputs operating at nine temperatures. All patterns are distinguishable from one another and reproducible from one experiment to the next. The topmost layer represents the winner-take-all competition (sensor 7 has the highest output for ammonia) and the bottommost layer represents the loser-take-all competition (sensor 5 has the lowest output for ammonia).

Mean and Median Thresholding of Chemical Senor Arrays

Mean and median thresholding techniques choose a single ordered point for referencing the entire array. In the mean thresholding case, this point corresponds to the average or mean value of the inputs in the array. In the median thresholding case, this reference point is the median value, where the half of the inputs above this ordered reference value are activated and the other half of the inputs are inactivated when generating a binary output pattern that is representative of chemicals in the sensing environment. Obvious differences exist between the mean and median thresholding processes. Mean thresholding tends to detect only the points of maximum response or interest in an array while median thresholding forces half of the points in an array to be of interest. As a result, mean thresholding is better suited to such applications as breath-alcohol monitoring, where only the primary chemical in the sensing environment (drinking alcohol) is of interest. Median thresholding, on the other hand, is better suited to applications where regions of chemical activity are of interest, rather than just the dominant chemical. Such applications include on-line monitoring of chemicals during such manufacturing processes as beer brewing and pharmaceutical drug production. The median thresholding elements presented in this chapter are also adjustable to ratios other than the 50/50 on/off ratio defined by the median. The bias currents may be easily adjusted to generate any ratio of on/off outputs. This capacity for adjustment is very useful for adapting the median thresholding elements to a variety of applications.

Rank-Order Filtering of Chemical Sensor Arrays

Inherently, rank-order filtering generates more information about an array of analog sensory inputs than the thresholding systems. A rank-order filtered response pattern, thresholds a sensory image based on J reference values, where J is the number of layers in the rank-order filtering process (winner, loser, runner-up, etc.). In contrast, each of the thresholding techniques, mean and median, only a single reference point for threshold the entire array of inputs. This multiple reference that is characteristic of the rank-order filtering circuits makes them useful for systems where a great deal of information about the sensing environment is required. Because the degree of data compression in the rank-order filtering is not as severe as in the thresholding process, this technique is also subject to more error and fluctuation than thresholding; resilience to error can be improved by using one of the preprocessing techniques described in Chapter 4 on a homogeneous cluster of sensors before rank-order filtering is performed. Because the rank-order filtering circuits can be adjusted to determine any number of ranks in an array, they may be used in a wide variety of applications, ranging from simple detection of carbon monoxide in the home to the on-line monitoring of manufacturing processes.

5.6 Comparison to Existing Systems

Other signal processing systems have been designed and fabricated in the research community to discriminate among the chemicals distinguished using techniques presented in this chapter. For the most part, these systems have required the use of a digital computer to make the final discrimination determination and in rare cases, contain sensors and processing elements that have been designed to be integrated onto a single substrate. Table 5.2 summarizes some of the research efforts in chemical discrimination microsystems in the past decade and cites some of the differences, advantages, and disadvantages of these systems over the signal processing architectures discussed in this chapter.

The primary advantage of the techniques presented in this chapter are their potential for use in low, low-cost systems (less than $20/unit) for chemical discrimination and their ease of use in large, parallel processed arrays of chemical sensors. This latter advantage will become more important as the state of chemical sensing microsystem technology becomes more advanced, requiring larger arrays of chemical sensors and more compact, multi-task, on-chip processing for efficient transfer of sensor information to more complex processing systems off-chip.

TABLE 5.2: Comparison of On-Chip Ordering Techniques to Other Chemical Analysis Systems

System	Discrimination Capability	Advantages	Disadvantages	REF
Mean Thresholding	Ammonia, Butanol, CO	low-cost, outputs independent of concentration and drift	Discrimination among only broad families of chemicals	
Median Thresholding	Ammonia, Butanol, CO	low-cost, outputs independent of concentration and drift	Discrimination among only broad families of chemicals	
Rank-Order Filtering	Ethanol, Isopropanol, CO	low-cost, outputs independent of concentration and drift, fine discrimination ability	Normalization of data reduces discrimination capability	
Fuzzy Neural Networks: software-trained, software-run	Ethanol, CO, Methanol	Single Output per Chemical	Requires computer, higher cost, outputs dependent on concentration	[59]
ASIC used to cross-correlate sensor signals	Butanol, Ethanol	Compact, low-cost sensing system	Outputs dependent on concentration	[60]
Neural Network: software-trained, hardware-run	Alcohols: ethanol, isopropanol	does not require microcomputer for operation	hardware can only accommodate up to 20 sensors	[40]