Author manuscript, published in Pattern Recognition Letters, vol. 32, no. 13, pp. 1535-1543, 2011

Detection of Soldering Defects in Printed Circuit Boards with Hierarchical Marked Point Processes

Csaba Benedek

Distributed Events Analysis Research Group, Computer and Automation Institute, Kende u. 13-17, Budapest, Hungary
Dept. of Electronics Technology, Budapest University of Technology and Economics, H-1111, Budapest, Goldmann tér 3

Abstract

In this paper we introduce a probabilistic approach for optical quality checking of Solder Pastes (SP) in Printed Circuit Boards (PCB). Dealing with unregistered image inputs, our task is to address at the same time SP identification, and detection of special soldering errors, called *scooping*. For this reason we introduce a novel Hierarchical Marked Point Process (HMPP) framework, which is able to handle the paste and scooping extraction problems simultaneously, so that the SPs and included scoops have a parent-child relationship. A global optimization process attempts to find the optimal configuration of entities, considering the observed data, prior knowledge, and interactions between the neighboring circuit elements. The proposed method is evaluated on a real PCB image set containing more than 3000 SPs and 600 scooping artifacts. A morphology-based baseline method is also introduced for the problem and used as reference for qualitative and quantitative comparison against the proposed model.

Key words: Marked point process, PCB, scooping

1. Introduction

Automatic optical inspection (AOI) is a crucial step in the manufacturing process of Printed Circuit Boards (PCBs). The quickly increasing resolution, quality and speed of industrial cameras have recently opened several new prospects in image-based verification. At finer scales a significant amount of information is revealed, which calls for shifting from simple segmentation (Comaniciu et al., 2002) or morphology-based (Soille and Vogt, 2009) investigations towards a hierarchical modeling approach of the PCB structure, focusing jointly on circuit regions, individual Circuit Elements (CEs), CE interactions and relevant CE parts.

This paper deals with the detection of a PCB printing defect, which derives also from a multiscale problem domain. Nowadays the most widespread assembling technology of electronic circuit modules applies reflow soldering (Krammer and Sinkovics, 2010). The amount of solder paste deposited affects the reliability and strength of the reflowed solder joint (Mannan et al., 1993). Solder paste scooping (see Fig. 1) occurs when the squeegee blade enters into the stencil aperture and removes paste from the center of the print. This is caused by high print pressure and can be exaggerated by rubber blades and large stencil apertures. A single scoop in a PCB does not cause a critical quality problem, however, if the number and summarized volume of such artifacts surpass a given threshold, the board should be withdrawn. Since in optical verifying systems PCBs are continuously moving on conveyor belts, the recognition algorithm must deal with unregistered image inputs and the circuit elements should be first located. Thus, the problem is hierarchical in the sense that we have to extract the solder pastes of the board and detect the included scoop within each imperfect paste simultaneously.

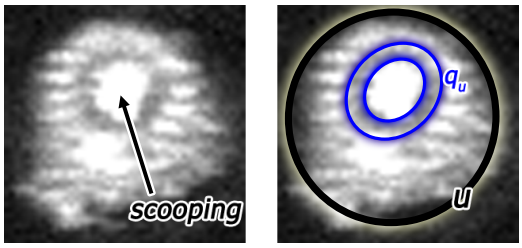

Fig. 1. Solder paste u with a scoop q_u ($10\mu m$ res. image)

Several previous AOI methods use mathematical morphology (Ionut, 2008; Soille and Vogt, 2009; Lin et al., 2009) for investigating geometric structures in binary images. Usually a series of morphological transformations is applied such as erosions, geodesic dilations, reconstruction by dilation, anchored skeletonisation, etc. The main weak point of these techniques is that they critically rely on the input binarization step performed by thresholding the grayscale input. However, estimating ap-

Document version of the MTA SZTAKI Publication Repository, http://eprints.sztaki.hu/

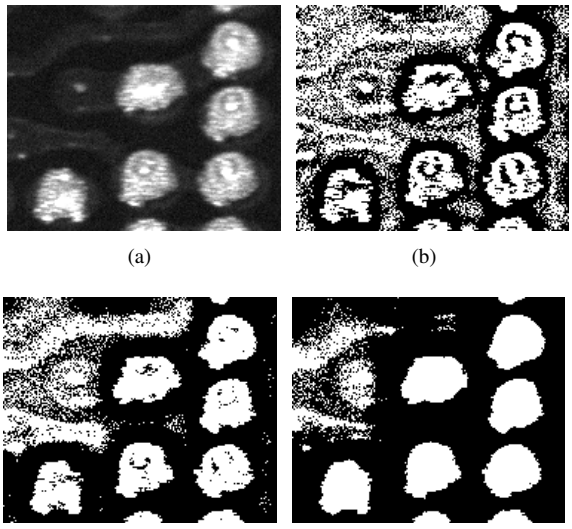

Fig. 2. Demonstration of Sauvola's locally adaptive thresholding algorithm (Sauvola and Pietikäinen, 2000) with different parameter settings

propriate global or local thresholds may be difficult, as well as, due to contrast defects and slight illumination variations, the separation of CEs from the background can be imperfect (Li et al., 2010). Even with applying efficient locally adaptive algorithms (Sauvola and Pietikäinen, 2000), fine structures within the CEs such as scooping are often hard to describe from a binarized image (see Fig. 2).

The quality control task is often interpreted as a change detection problem, comparing the examined PCB to a faultless reference board (Tsai and Lin, 2003; Tsai et al., 2003) or match the CEs to template entities (Tsai and Yang, 2005). However, in the addressed circuit images, even the perfect solder pastes are highly inhomogeneous, and contain irregularly distributed lighter and darker regions. This phenomenon may make the comparison of pixel intensities or textures of the extracted pastes misleading. Meanwhile, the scoops exhibit various size, orientation and contrast parameters (see Fig. 10-12), which can corrupt template matching approaches.

Recently, Marked Point Processes (MPP) (Descombes and Zerubia, 2002; Descombes et al., 2009) have become popular in 2D and 3D pattern recognition, since they can efficiently model a population of an unknown number of entities and consider the geometry of objects. Following the inverse modeling approach, a fitness function is defined on the space of the possible object configurations, and an optimization process attempts to find the configuration with the highest confidence. To keep the computational the complexity tractable, stochastic optimization techniques are usually adopted, which consist of object generation (*birth*), removal (*death*) and perturbation steps. In such models, validation of objects proposed by the *birth* process is a key point, as it is used to decide if a candidate should be kept, deleted or modified.

Although conventional MPP models are well established tools in various pattern recognition application areas, they do not handle object hierarchy, thus they are inappropriate to model complex multi-level scenarios. Stepping forward, we describe

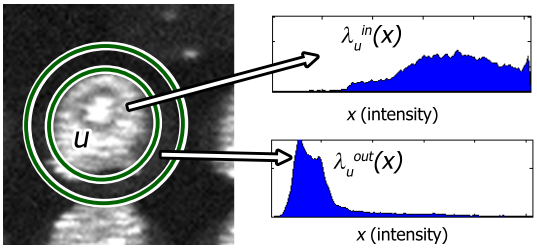

Fig. 3. Utility of the $\lambda_u^{\text{in}}(x)$ and $\lambda_u^{\text{in}}(x)$ histograms for the parent data term calculation

here the hierarchy between solder paste objects and internal scoops as a parent-child relationship embedded into a novel Hierarchical MPP (HMPP) framework. The appearance of a child object is affected by its parent entity, considering geometrical and spectral constraints, such us a parent object contains the child, or the contrast within the parent object influences the intensity characteristic of the child entity. Besides, prior knowledge is exploited in modeling the interaction between different objects.

2. Problem definition

The input of the proposed framework is an image taken from a PCB, having a 2D pixel lattice S. The goal is to extract and separate the Solder Pastes (SP) in the PCBs, while detecting and describing the Scoop Objects (SO) appearing in some of the pastes (Fig. 1). We model a SP by an ellipse; and a SO by two concentric ellipses, as it consists of a bright central region (inside the internal ellipse) and a relatively darker elliptical ring (region between the internal and external ellipses) enclosed by the brighter SP patch (Fig. 1).

A given SP ellipse, denoted by u, is described by five geometric parameters: x and y center coordinates in the S lattice, semimajor resp. semiminor axes, and the orientation angle. In addition, u may contain a SO, q_u, which is determined by its center pixel, semi-axes of the internal and external ellipses, respectively, and the scoop orientation. We consider in the following q_u as a child object assigned to each u, and denote by $q_u = \text{nil}$ if paste u contains no scoop. We denote by $\mathcal{H}$ the set of all the possible (u, q_u) parent-child object structures.

3. Feature selection

In this section we introduce image based features which can characterize the SP and SO candidates. Based on the extracted descriptors, we define an energy function $\varphi_d(u) : \mathcal{H} \to [-1, 1]$, which evaluates an object hypothesis considering local image data, but disregarding other detected objects on the board. $\varphi_d(u)$ is interpreted as a negative fitness value of an arbitrary (u, q_u) pair. By definition, a SP with $\varphi_d(u) < 0$ is called an attractive object, and we aim to construct the $\varphi_d(u)$ function so that attractive objects correspond to the true circuit entities.

The data term $\varphi_d(u)$ is decomposed into a parent term $\varphi_d^p(u)$ and a child term $\varphi_d^c(u, q_u)$. In the PCB images the *parent* SPs can be modeled as bright ellipses surrounded by darker back-

Document version of the MTA SZTAKI Publication Repository, http://eprints.sztaki.hu/

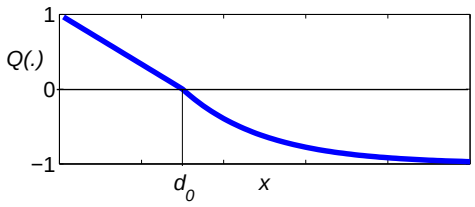

Fig. 4. Plot of the $Q(x, d_0)$ function

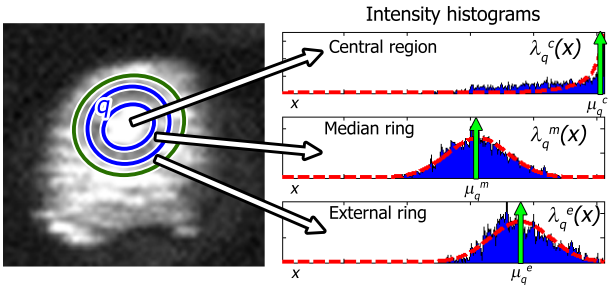

Fig. 5. Child data term calculation

ground. To evaluate the contrast between the pastes and the board, we calculate the Bhattacharya (Descombes et al., 2009) distance $d_B(u)$ between the pixel intensity distributions of the internal SP regions and their boundaries:

$$d_B(u) = 1 - \int \sqrt{\lambda_u^{\text{in}}(x) \cdot \lambda_u^{\text{out}}(x)} dx$$

where $\lambda_u^{\text{in}}(x)$ (resp. $\lambda_u^{\text{out}}(x)$) is the empirical gray level distribution of pixels belonging to u (resp. a concentric elliptical ring around u), as shown in Fig. 3.

In the next step, we construct the parent energy subterm $\varphi_d^p(u)$, so that we attempt to satisfy $\varphi_d^p(u) < 0$ for real objects and $\varphi_d^p(u) \geq 0$ for false candidates. For this purpose, we project the $d_B(u)$ feature domain with a monotonously decreasing function (see also Fig. 4):

$$\varphi_d^p(u) = Q(d_B(u), d_0) = \\ = \begin{cases} \left(1 - \dfrac{d_B(u)}{d_0}\right) & \text{if } \; d_B(u) < d_0 \\ \exp\left(-\dfrac{d_B(u) - d_0}{0.1}\right) - 1 & \text{if } \; d_B(u) \geq d_0 \end{cases} \quad (1)$$

where d_0 is the object-acceptance threshold.

The construction of the *child's data term* $\varphi_d^c(u, q_u)$ is based on similar principles. We use $\varphi_d^c(u, \text{nil}) = 0$, otherwise we distinguish three regions of each scoop: the central bright ellipse, the darker median ring and the bright external ring, as shown in Fig. 5. Experimental evidence proves that for a real scoop q, the gray level histogram of the central region, $\lambda_q^{\text{c}}(x)$ follows a skewed distribution, while the medium and external region histograms ($\lambda_q^{\text{m}}(x)$ resp. $\lambda_q^{\text{e}}(x)$) can be approximated by Gaussian densities. Let us denote by μ_q^{c} the peak location of $\lambda_q^{\text{c}}(x)$, and by μ_q^{m} resp. μ_q^{e} the empirical mean values of the $\lambda_q^{\text{m}}(x)$ resp. $\lambda_q^{\text{e}}(x)$ distributions.

We characterize a scoop candidate by the following three features:

(i) μ_q^{c}: brightness of the central region

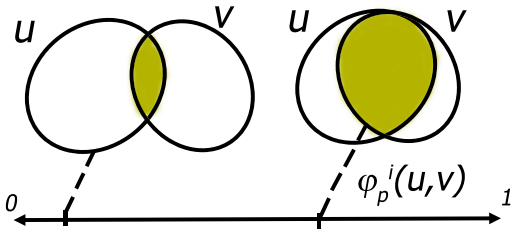

(a) Intersection area at parent level

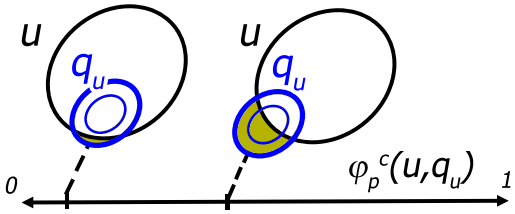

(b) Scoops overhanging area w.r.t the parent SP object

Fig. 6. Calculation of the prior terms

(ii) intensity ratio $\mu_{q_u}^{\text{c}}/\mu_{q_u}^{\text{m}}$: contrast between the central region and median ring

(iii) intensity ratio $\mu_{q_u}^{\text{e}}/\mu_{q_u}^{\text{m}}$: contrast between the external ring and median ring

We assign to each feature an energy term using the Q-function similarly to (1).

To enforce the simultaneous appropriateness of the (i)-(iii) features, the child's data-energy value is calculated using the maximum operator (logical AND in the negative fitness domain) from the subterms of the three constraints:

$$\varphi_d^c(u, q_u) = \max\Big(Q(\mu_{q_u}^{\text{c}}, d^c), \\ Q(\mu_{q_u}^{\text{c}}/\mu_{q_u}^{\text{m}}, d^{\text{cm}}), \\ Q(\mu_{q_u}^{\text{e}}/\mu_{q_u}^{\text{m}}, d^{\text{em}})\Big) \quad (2)$$

Free parameters of the scoop intensity model are d^c, d^{cm} and d^{em}, which have been set based on manually evaluated training data. Maximum Likelihood Estimators (MLE) can optimize the parameter values as detailed in (Chatelain et al., 2009).

On the other hand, the width of object boundaries for calculating the $d_B(u)$ and $\lambda_q^{\text{e}}(.)$ features are determined using prior size information about the expected entities.

Finally, the complete data term of the SP candidate u is obtained as

$$\varphi_d(u) = \varphi_d^p(u) + \varphi_d^c(u, q_u)$$

4. Hierarchical Marked Point Process Model

After investigating independent circuit elements (Sec. 3), we characterize complete object configurations in this section. We model the PCB image as composition of an arbitrary number of SPs whose positions and geometric parameters are realizations of a MPP (Descombes and Zerubia, 2002). As a novelty of our approach, the scoops realize an embedded MPP within the SPs. Let us denote by Ω the space of all configurations of a finite number of SP objects. We will refer to a given object configuration in Ω by ω, where $\omega = \emptyset$ or $\omega = \{u_1, \ldots, u_n\}$ for

Document version of the MTA SZTAKI Publication Repository, http://eprints.sztaki.hu/

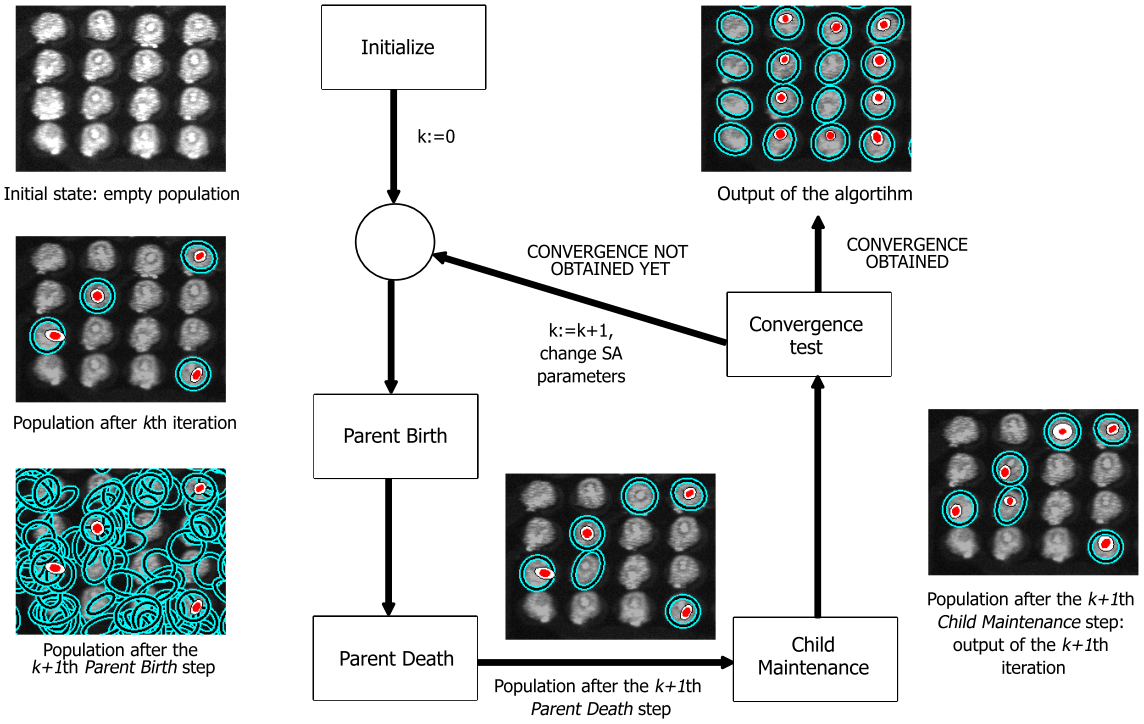

Fig. 7. Demonstration of the Hierarchical MBD optimization algorithm. Comparing the populations after the *k*th and *k+1*th iterations in the example images, we can observe that the applied *Parent Birth-Death* step pair recognized two new SP objects, while the Child Maintenance process detected their corresponding scoops and also perturbated three already existing SOs.

an arbitrary positive integer n. We also define a neighborhood relation in the SP object space, $u \sim v$ if their ellipses intersect.

Next, we introduce a $\Phi(\omega)$ energy function on the Ω configuration space, so that we aim to fulfill that efficient configurations exhibit relatively lower energies. The energy function takes into account the way the entities fit to the optical data (the data energy $\Phi_d(\omega)$) and geometric interactions between contacting objects (the prior energy $\Phi_p(\omega)$):

$$\Phi(\omega) = \kappa \cdot \Phi_d(\omega) + \Phi_p(\omega) \tag{3}$$

where κ is a positive weighting factor between the two energy parts, we used a constant $\kappa = 2$.

Based on the $\varphi_d(u)$ object-energy values defined in Sec. 3, the data term of the configuration is derived as:

$$\Phi_d(\omega) = \sum_{u \in \omega} \varphi_d(u)$$

We continue the description with the $\Phi_p(\omega)$ energy part, which prescribes *prior geometric* constraints in the model.

Since we aim to extract individual SP entities, we must penalize overlapping between different SP ellipses. Thus, the energy term of parent-level interactions penalizes object pairs according to the $\varphi_p^i(u, v)$ normalized overlapping area as shown in Fig. 6(a).

On the other hand, we should also penalize if a scoop q_u overhangs its parent paste u: this is measured by the $\varphi_p^c(u, q_u)$ overhanging area normalized by the area of the scoop (see Fig. 6(b)). By definition, we use $\varphi_p^c(u, \mathrm{nil}) = 0$.

The prior energy of the population is calculated as the sum of the overlapping and overhanging energy terms over all objects and object pairs:

$$\Phi_p(\omega) = \sum_{\substack{u,v \in \omega \\ u \sim v}} \varphi_p^i(u, v) + \sum_{u \in \omega} \varphi_p^c(u, q_u)$$

As the energy function $\Phi(\omega)$ is completely defined, the optimal SP population $\hat{\omega}$ is obtained as the Maximum Likelihood (ML) configuration estimate:

$$\hat{\omega} = \operatorname*{argmin}_{\omega \in \Omega} \Phi(\omega). \tag{4}$$

5. Optimization and computational requirements

For the estimation of the optimal object configuration (4) we have developed a hierarchical extension of the recent Multiple Birth and Death (MBD) method (Descombes et al., 2009). The flowchart of the relaxation algorithm is demonstrated in Fig. 7, and the steps of the process are detailed in Fig. 8.

As the optimization method is stochastic, the objects are generated in the birth step with random size parameters (major resp. minor axes of SP and SO ellipses), following size distributions estimated from sample objects. In case of different CE types present on the board (see Fig. 11), the prior distributions are multimodal: we use a Mixture of Gaussians approximation of the empirical size histograms for this purpose. Regarding the *relaxation* parameters, we followed the guidelines provided in (Descombes et al., 2009) and used $\delta_0 = 10000$, $\beta_0 = 20$ and geometric cooling factors $1/0.96$.

An important issue is to address the computational cost of the algorithm in terms of industrial requirements. We performed our tests on a standard desktop computer, where processing

Document version of the MTA SZTAKI Publication Repository, http://eprints.sztaki.hu/

Algorithm: Hierarchical MBD optimization

Step 1: Initialization

(a) Start with an empty population ($\omega = \emptyset$).

(b) Set initial Simulated Annealing parameters inverse temperature $\beta = \beta_0$ and discretization step $\delta = \delta_0$.

(c) Let b_0 be a basic birth frequency.

Step 2: Parent Birth

For each pixel s of lattice S call the following *SP Generation* procedure with a probability $\delta \cdot b_0$:

(i) Create a new SP object u with center s.

(ii) Set u's axis length parameters randomly following prior size distributions, and a random orientation.

(iii) Use null-scoop initially: $q_u = \mathrm{nil}$.

(iv) Add u to the current configuration ω.

Step 3: Parent Death

Consider the current configuration ω:

(a) Create a list of the $u \in \omega$ objects sorted from the highest to the lowest $\varphi_d^p(u)$ values.

(b) For each object u taken in this order, compute the cost of deleting u from ω w.r.t. the global configuration energy:

$$\Delta\Phi_\omega^u = \Phi(\omega/\{u\}) - \Phi(\omega)$$

(c) Derive the $p_d(u)$ *death rate* as follows:

$$p_d(u) = f\left(\Delta\Phi_\omega^u\right) = \frac{\delta \cdot \exp\left(-\beta \cdot \Delta\Phi_\omega^u\right)}{1 + \delta \cdot \exp\left(-\beta \cdot \Delta\Phi_\omega^u\right)} \tag{5}$$

(d) Remove u from ω with a probability $p_d(u)$.

Step 4: Child Maintenance

For each u object in ω:

(a) Generate a new scoop candidate q'_u which is either $q'_u = \mathrm{nil}$; or we pick up a random point covered by the ellipse of u, and appoint it as the center of q'_u, while orientation and axes parameters of q'_u are set randomly.

(b) Calculate the energy cost of exchanging q_u to q'_u as follows:

$$\Delta\varphi(u, q_u, q'_u) = \varphi_p^c(u, q'_u) + \varphi_d^c(u, q'_u) - \left(\varphi_d^c(u, q_u) + \varphi_p^c(u, q_u)\right)$$

(c) Calculate the *scoop exchange likelihood* using the $f(.)$ function defined by (5):

$$p_c(u) = f\left(\Delta\varphi(u, q_u, q'_u)\right)$$

(d) Replace q_u with q'_u with a probability $p_c(u)$.

Step 5: Convergence Test

If the process has not converged yet, increase the inverse temperature β and decrease the discretization step δ with a geometric scheme, and **GOTO Step 2.**, otherwise **STOP**. The convergence is obtained when all the SP objects added during the birth step, and only these ones, have been killed during the death step, meanwhile the *Child Maintenance* does not report any more changes.

Fig. 8. Pseudo code of the developed Hierarchical Multiple Birth and Death optimization algorithm

1MP input images with 76 SP entities took around 10-20 sec. Apart from further source code optimization, there are several options to accelerate the method from algorithmic points of view. Firstly, similarly to other graph-like model structures (such as Markov Random Fields (Szirányi et al., 2000)), MPP optimization algorithms can be parallelized and adopted for multiprocessor architectures (Bambos and Chen, 1994). Secondly, in case of input images with sufficiently high Signal to Noise Ratio (SNR), the convergence of the MBD relaxation process can be significantly sped up with using so called "birth maps" as shown in (Descombes et al., 2009).

Another real-life application constraint is that the industrial implementation of the method should have a deterministic processing speed. Experimental studies confirm that the MBD algorithm with constant geometric cooling schedule shows a well predictable convergence, and the required number of iterations can be estimated a priori. Thus practically the process can be stopped after a fixed number of steps, without considerable degradation of the output quality.

On the other hand, our test data providers confirmed that with their current technology, the stencil printing process takes around 15 second for a four-piece batch of PCBs. This means that even with the hardware-software configuration used in our tests, we can nearly meet the real time verification requirement if we set four simultaneously working processing units.

6. Further circuit error investigations using the HMPP method

Although the proposed method has been developed for the scooping detection task, various artifacts may also be investigated in a similar way using the introduced HMPP model. For

Document version of the MTA SZTAKI Publication Repository, http://eprints.sztaki.hu/

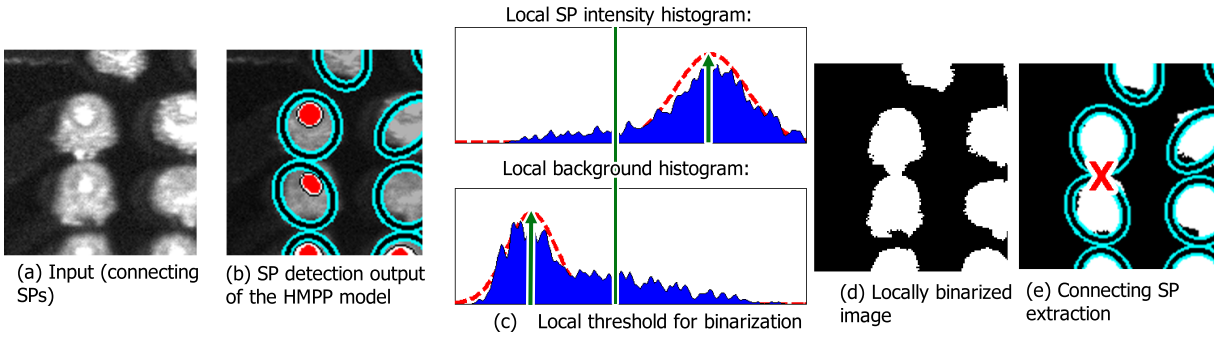

Fig. 9. **Short circuit detection**: demonstration of the consecutive steps of the connecting SP extraction algorithm. Local threshold for binarization is estimated based on the output of the HMPP detector. Thereafter, based on the locally binarized mask and the obtained SP ellipses, connecting elements can be detected by morphological operations.

example, critical PCB errors are the short circuits caused by thin lines connecting the neighboring SPs. To deal with this problem, our model can be extended in a straightforward way (Ortner et al., 2008): we can use in parallel line segment objects to detect the short circuit connections and ellipses modeling the SPs. For the line segments, similar data and interaction terms can be constructed to the cases of the SPs. The output entity configuration, which consists of an unknown number of ellipses and line segments, can be obtained with an analogous optimization process.

As well, with a few postprocessing operations, we can detect if neighboring SPs touch each other as in Fig. 9 and 11(a). The proposed steps are as follows:

– For each neighboring SP-pair detected by HMPP (see Fig. 9(b)), we calculate the intensity histograms of the paste regions (ie. inside the ellipses) and the background regions (ie. around the two ellipses) as shown in Fig. 9(c).
– Based on these local measurements, we separate the paste and board areas by thresholding the image part with the mean of the two filtered histogram peaks (Fig. 9(c)-(d)).
– In the considered binary mask we check with morphological operators if the two SP blobs are connected, so that we start a floodfill propagation from the center of first SP and measure if the resulting floodfill mask covers the estimated region of the second SP ellipse with more than 80% (Fig. 9(e)).

Result of touching element detection is also demonstrated in Fig. 11(e). We can observe that the connecting CEs are correctly separated by the proposed HMPP model, but the short circuit errors are detected at the same time.

7. Reference method

Although the same automated scoop extraction task has not been addressed in the literature of pattern recognition yet, it is relevant to compare the performance of the proposed HMPP model to more conventional approaches. As noted in the introduction, most previous template-free AOI methods were based on morphological investigations (Ionut, 2008; Soille and Vogt, 2009; Lin et al., 2009). Therefore, based on tools used in the earlier techniques, we have also developed a morphology-based solution (*Morph*) and used it as a baseline against our model.

The *Morph* technique exploits that in general SPs can be identified as relatively light but non-homogeneous blobs in the boards, meanwhile scoop seeds correspond to the lightest compact parts of the SPs, surrounded by a darker ring. Therefore two thresholding operations are applied on the input image: the first one uses a lower threshold value τ_l, and results in the binary SP candidate mask Γ_{SP}. The second threshold τ_h enables us to extract the lightest image parts only, so that we attempt to realize that in a second mask binary image, Γ_{SO}, the *white* regions correspond to the scoop center areas. Since light SP parts also occur independently of scooping, a verification process is needed, which aims to remove false SO candidates. This post-processing step must also ensure that each SP contains one scoop at most, which is a consequence of the manufacturing process.

By testing different threshold configurations for the scoop extraction step, we concluded the following:

– Separation of the central region from the median ring proved to be often reasonable with an appropriate threshold, however, several false candidates also appeared (see Fig. 10(f))
– Separation of the median ring and the external region failed with thresholding techniques, because the intensity domains were significantly overlapping.

Therefore, instead of the three-region scoop model of HMPP (Fig. 5), we characterized a SO in the *Morph* method by a two-region model which consists of the center-ellipse and its neighborhood.

In details, the *Morph* method contains the following main steps, which are also demonstrated in Fig. 10:

SP separation and detection:

1 Calculate τ_l by Otsu's automated threshold estimation method (Otsu, 1979)
2 Obtain the SP candidate mask Γ_{SP} by thresholding the input image with τ_l (see Fig. 10(b))
3 Identify and separate SPs by Connected Component Analysis (CCA1) on Γ_{SP} (Fig. 10(c))
4 Estimate the shape of each extracted SP by fitting an ellipse to the binary mask of the paste (Fig. 10(d))

SO detection

5 Calculate τ_h as the 90% of the maximal gray level value

Document version of the MTA SZTAKI Publication Repository, http://eprints.sztaki.hu/

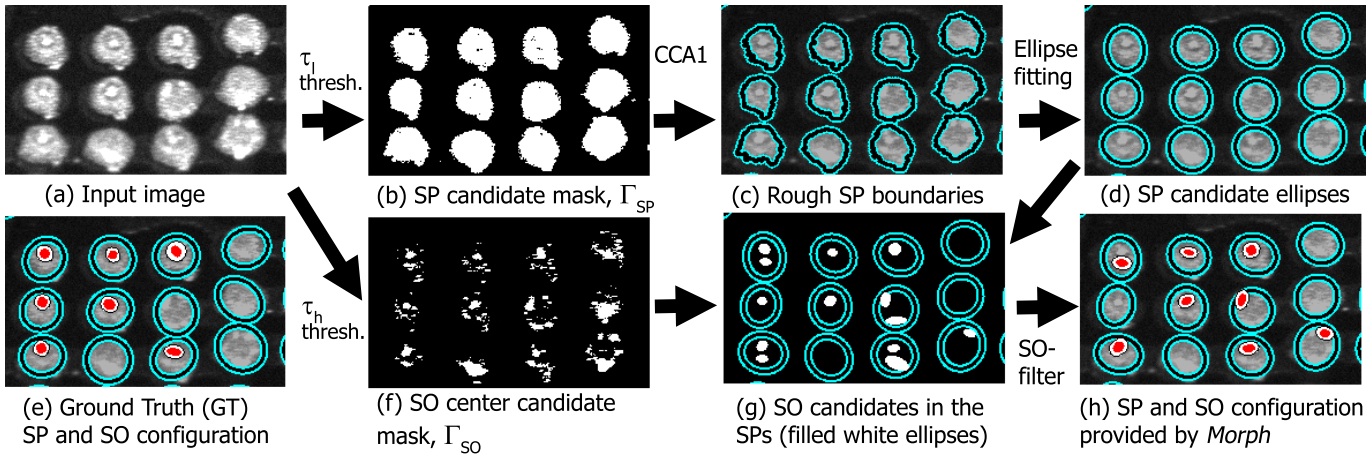

Fig. 10. **Reference method**: demonstration of the consecutive steps of the *Morph* algorithm

in the image [1]

6 Obtain SO center candidate mask $\Gamma_{\rm SO}$ by thresholding the input image with τ_h (Fig. 10(f)). Denote by $\gamma_{\rm SO}(s) \in \{0,1\}$ the class value of pixel s in $\Gamma_{\rm SO}$, where 1 corresponds to the *white* regions.

7 Within each SP ellipse u detected in Step 4, select the appropriately sized blobs of $\Gamma_{\rm SO}$, which may correspond to SP center regions (CCA2), and approximate them with ellipses: $b_1(u), \dots b_{k_u}(u)$ (Fig. 10(g)). Denote by $s \in b_i(u)$ if pixel s corresponds to the scoop candidate $b_i(u)$.

8 Compute a fitness value for each $b_i(u)$ $i = 1 \dots k_u$, $\forall u$, so that we prescribe for a relevant scoop candidate that the SO seed ellipse should consist of "white" pixels in $\Gamma_{\rm SO}$ in majority, while the neighborhood of the seed should be "black":

– Calculate an internal filling factor $f_{\rm int}$ as:

$$f_{\rm int}(u,i) = \frac{\#\{s : s \in b_i(u), \gamma_{\rm SO}(s) = 1\}}{\#\{s : s \in b_i(u)\}}$$

– Denote by $\widehat{b}_i(u)$ a 4 pixel wide elliptical ring around $b_i(u)$, and calculate an external filling factor $f_{\rm ext}$ as:

$$f_{\rm ext}(u,i) = \frac{\#\{s : s \in \widehat{b}_i(u), \gamma_{\rm SO}(s) = 0\}}{\#\{s : s \in \widehat{b}_i(u)\}}$$

– Finally, derive the fitness value of the scoop candidate $b_i(u)$ as the average of the internal and external factors:

$$f(u,i) = (f_{\rm int}(u,i) + f_{\rm ext}(u,i))/2$$

9 Find to each SP ellipse u the index of its most likely scoop-candidate $i_{\max}(u) \in \{1, \dots, k_u\}$:

$$i_{\max}(u) = \operatorname*{argmax}_i f(u,i)$$

10 If the $f(u, i_{\max})$ fitness value surpasses an acceptance threshold, assign a SO to u with seed ellipse $b_{i_{\max}}(u)$ (Fig. 10(h)). Otherwise use nil scoop for u.

[1] At this step, we have also tested the slower locally adaptive thresholding procedure, but due to the large number of oversaturated regions in the test images, we could not obtain considerable improvement by that modification.

Note that the *Morph* method has a few free parameters, such as size range for SP resp. SO centers, and the SO acceptance threshold. We optimized these parameter in our test manually based on training samples, since we were primarily interested in the fundamental limitations of the approach, and its comparison to the proposed HMPP model.

8. Evaluation

We have tested our method on various real PCB images (see Fig. 13). The largest available test set has been provided by Nokia, which consists of 10μm resolution photos of 40 different circuits, each of them contains 76 SPs. In aggregate 613 SOs are present in the 3040(=40·76) solder pastes of the test set. The quality of the different PCBs shows a large variety: 18 out of the 40 analyzed images contain less than 10 scoops/board, 14 circuits include 10-30 scoops, and each of the remaining 8 ones holds more than 30 SOs. Meanwhile, the test images exhibit significantly different and spatially varying contrast due to the quick motion of the production line and the optical limitations of the industrial cameras. For comparison, both the *Morph* and the proposed *HMPP* approaches were evaluated on the whole Nokia data set. Images from the other considered circuits do not contain scooping, however, the conditions for CE-background separation are more challenging (see Fig. 13).

For accurate Ground Truth (GT) generation, we have developed an accessory program with graphical user interface, which enables us to add SP ellipses and SO ellipse pairs (i.e. objects) to the configuration manually, and the result appears immediately over the input image. The objects can be moved and the parameters (position, axes and angles) can be arbitrarily perturbated with pixel accuracy, using buttons and sliders. If needed, a zooming tool helps the accurate visual validation of the drawn objects. Thereafter, the GT objects (both for the training and test samples) were edited by members of our department and validated by experts of the manufacturing process.

We have experienced that the parameter estimation process of the parent SP extraction can be efficiently automated without manual parameter tuning, so that we apply a coarse foreground-

Document version of the MTA SZTAKI Publication Repository, http://eprints.sztaki.hu/

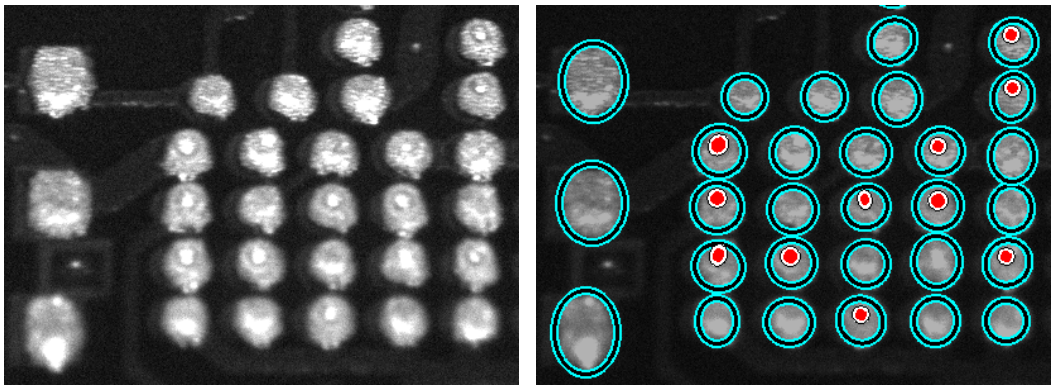

(a) Input image (b) Ground Truth

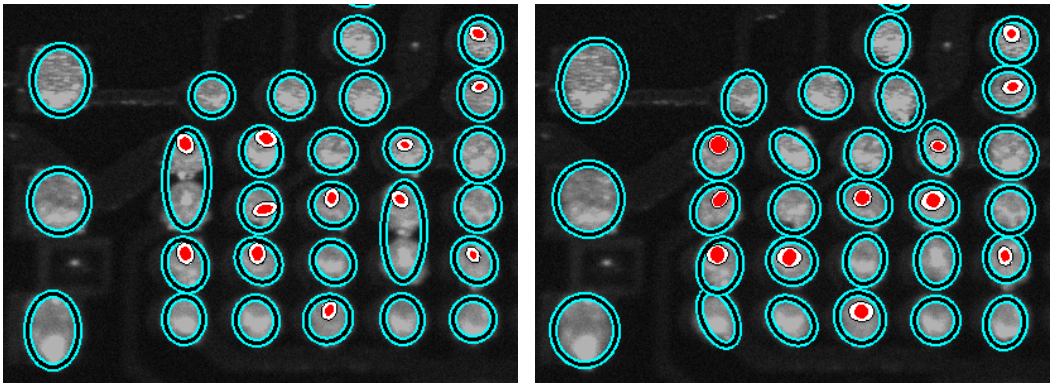

(c) Morph result (d) HMPP result

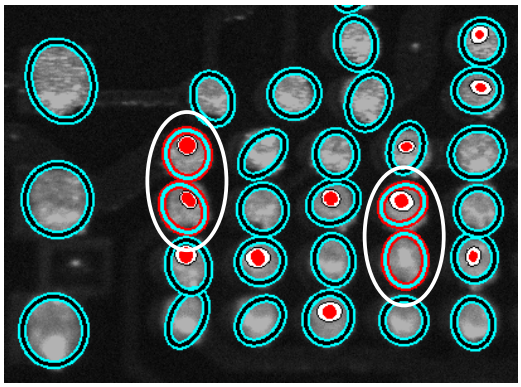

(e) HMPP + short circuit detection

Fig. 11. Results in a sample image with the *Morph* and the proposed *HMPP* techniques. Errors of the *Morph* method can be observed both at SP and at SO levels. In (e) we demonstrate the output of the short circuit detection algorithm applied after *HMPP* - connecting circuit element pairs are correctly separated first, then detected as short circuit errors

background separation by Otsu's thresholding method, then we set the acceptance threshold parameter d_0 based on statistical description of the coarsely identified SP and background regions. For tuning the SO intensity model we have relied on a supervised method, where in aggregate 15 manually selected scoops were used for training from three different PCB images of the Nokia set, and the remaining GT samples were only utilized for validating the detection results.

Some qualitative results of the hierarchical SP-SO extraction are shown in Fig. 11, 12 and 13. We have fulfilled the quantitative evaluation both for the SP and SO objects in the Nokia set.

Table 1
Quantitative evaluation results of joint solder paste & scoop extraction with the proposed model for the Nokia set. The used test data contains 3040 SP entities and 613 SOs.

	SP level		SO level		
Method	Rc	Pc	Rc	Pc	FM
Morph	98.6%	100%	78.8%	65.2%	71.4%
Prop. HMPP	99.5%	99.9%	90.7%	91.6%	91.1%

At both levels, we have counted the number of True Positive (TP), False Negative (FN) and False Positive (FP) objects in

Document version of the MTA SZTAKI Publication Repository, http://eprints.sztaki.hu/

the detection results. Thereafter, we calculated the Recall (Rc) and Precision (Pc) rates as:

$$\mathrm{Rc} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}} \qquad \mathrm{Pc} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FP}}$$

Finally, the detection was characterized by the *F-measure* (FM), which combines *Recall* and *Precision* in a single efficiency measure (Rijsbergen, 1979), by computing the harmonic mean of Rc and Pr:

$$\mathrm{FM} = \frac{2 \cdot \mathrm{Rc} \cdot \mathrm{Pc}}{\mathrm{Rc} + \mathrm{Pc}}$$

Numerical evaluation results are shown in Table 1. At SP level the detection was nearly perfect, minor errors are caused by notably blurred image regions and noisy patches in some of the images, where the proposed approach proved to be more robust than the *Morph* model (see also Fig. 11).

At SO level, differences between the two methods are significant, as the proposed HMPP model outperforms the *Morph* technique by nearly 20% regarding the FM rate. The main weak point of *Morph* proved to be the high false detection rate (see Pc values). This failure parameter can be decreased by adopting a more strict verification step, however, we observed that such a modification results in a drastic decrease of the Recall rate too. On the other hand, circuit technologists have confirmed that, based on our reported accuracy with the HMPP model, a fair statistical analysis of the scooping effects can be performed for quality characterization of the PCBs.

9. Conclusion

We have proposed a novel Hierarchical Marked Point Process (HMPP) framework adopted to optical scooping analysis in PCBs. The method incorporates Solder Paste extraction and Scooping error detection in a joint probabilistic approach. To demonstrate the limitations of previous PCB checking methods concerning the selected task, we have developed a reference technique, called *Morph*, from a sequence of conventional binary morphological operations. We have performed a detailed quantitative and qualitative comparative evaluation between the *Morph* and the *HMPP* models in a real high resolution PCB image set. Experiments confirmed the superiority of the proposed HMPP model and its usability for forthcoming industrial inspection systems.

Acknowledgment

The work of the author was partially supported by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences.

The author would like to thank László Jakab and Olivér Krammer for provision of the input images and for their kind remarks and advice regarding the technical aspects of the problem and task definition.

References

Bambos, N. and Chen, S. C. (1994). Optimality aspects of greedy schemes in parallel processing of random graph-structured jobs. *Probability in the Engineering and Informational Sciences*, 8:229–243.

Chatelain, F., Descombes, X., and Zerubia, J. (2009). Parameter estimation for marked point processes. application to object extraction from remote sensing images. In *Proc. Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR)*, Bonn, Germany.

Comaniciu, D., Meer, P., and Member, S. (2002). Mean shift: A robust approach toward feature space analysis. *IEEE Trans. on Pattern Anal. and Machine Intelligence*, 24:603–619.

Descombes, X., Minlos, R., and Zhizhina, E. (2009). Object extraction using a stochastic birth-and-death dynamics in continuum. *J. Math. Imaging and Vision*, 33:347–359.

Descombes, X. and Zerubia, J. (2002). Marked point processes in image analysis. *IEEE Signal Processing Magazine*, 19(5):77–84.

Ionut, B. (2008). Using mathematical morphology to detect the imperfections of the printed circuit boards. *Journal of Applied Computer Science & Mathematics*, 2(3):9–14.

Krammer, O. and Sinkovics, B. (2010). Improved method for determining the shear strength of chip component solder joints. *Microelectronics Reliability*, 50(2):235 – 241.

Li, Z., Yang, J., Liu, C., Liu, G., and Cheng, Y. (2010). Unsupervised range-constrained thresholding. *Pattern Recogn. Lett.* In press.

Lin, L.-Z., Zhou, L.-S., Wan, J.-D., and Qian, Z.-Q. (2009). Study of PCB automatic optical inspection system based on mathematical morphology. In *International Conference on Computer Technology and Development*, volume 2, pages 405–408, Los Alamitos, CA, USA. IEEE Computer Society.

Mannan, S., Ekere, N., Lo, E., and Ismail, I. (1993). Predicting scooping and skipping in solder paste printing for reflow soldering of smt devices. *Soldering & Surface Mount Technology*, 5(3):14 – 17.

Ortner, M., Descombes, X., and Zerubia, J. (2008). A marked point process of rectangles and segments for automatic analysis of digital elevation models. *IEEE Trans. Pattern Analysis and Machine Intelligence*.

Otsu, N. (1979). A threshold selection method from gray-level histogram. *IEEE Trans. on Systems, Man, and Cybernetics*, 4(2):62–66.

Rijsbergen, C. J. V. (1979). *Information Retrieval*. Butterworths, London, second edition.

Sauvola, J. and Pietikäinen, M. (2000). Adaptive document image binarization. *Pattern Recognition*, 33:225–236.

Soille, P. and Vogt, P. (2009). Morphological segmentation of binary patterns. *Pattern Recogn. Lett.*, 30(4):456–459.

Szirányi, T., Zerubia, J., Czúni, L., Geldreich, D., and Kato, Z. (2000). Image segmentation using markov random field model in fully parallel cellular network architectures. *Real-Time Imaging*, 6(3):195–211.

Tsai, D.-M. and Lin, C.-T. (2003). Fast normalized cross correlation for defect detection. *Pattern Recogn. Lett.*, 24(15):2625–2631.

Tsai, D.-M., Lin, C.-T., and Chen, J.-F. (2003). The evaluation of normalized cross correlations for defect detection. *Pattern Recogn. Lett.*, 24(15):2525–2535.

Tsai, D.-M. and Yang, C.-H. (2005). A quantile-quantile plot based pattern matching for defect detection. *Pattern Recogn. Lett.*, 26(13):1948–1962.

Document version of the MTA SZTAKI Publication Repository, http://eprints.sztaki.hu/

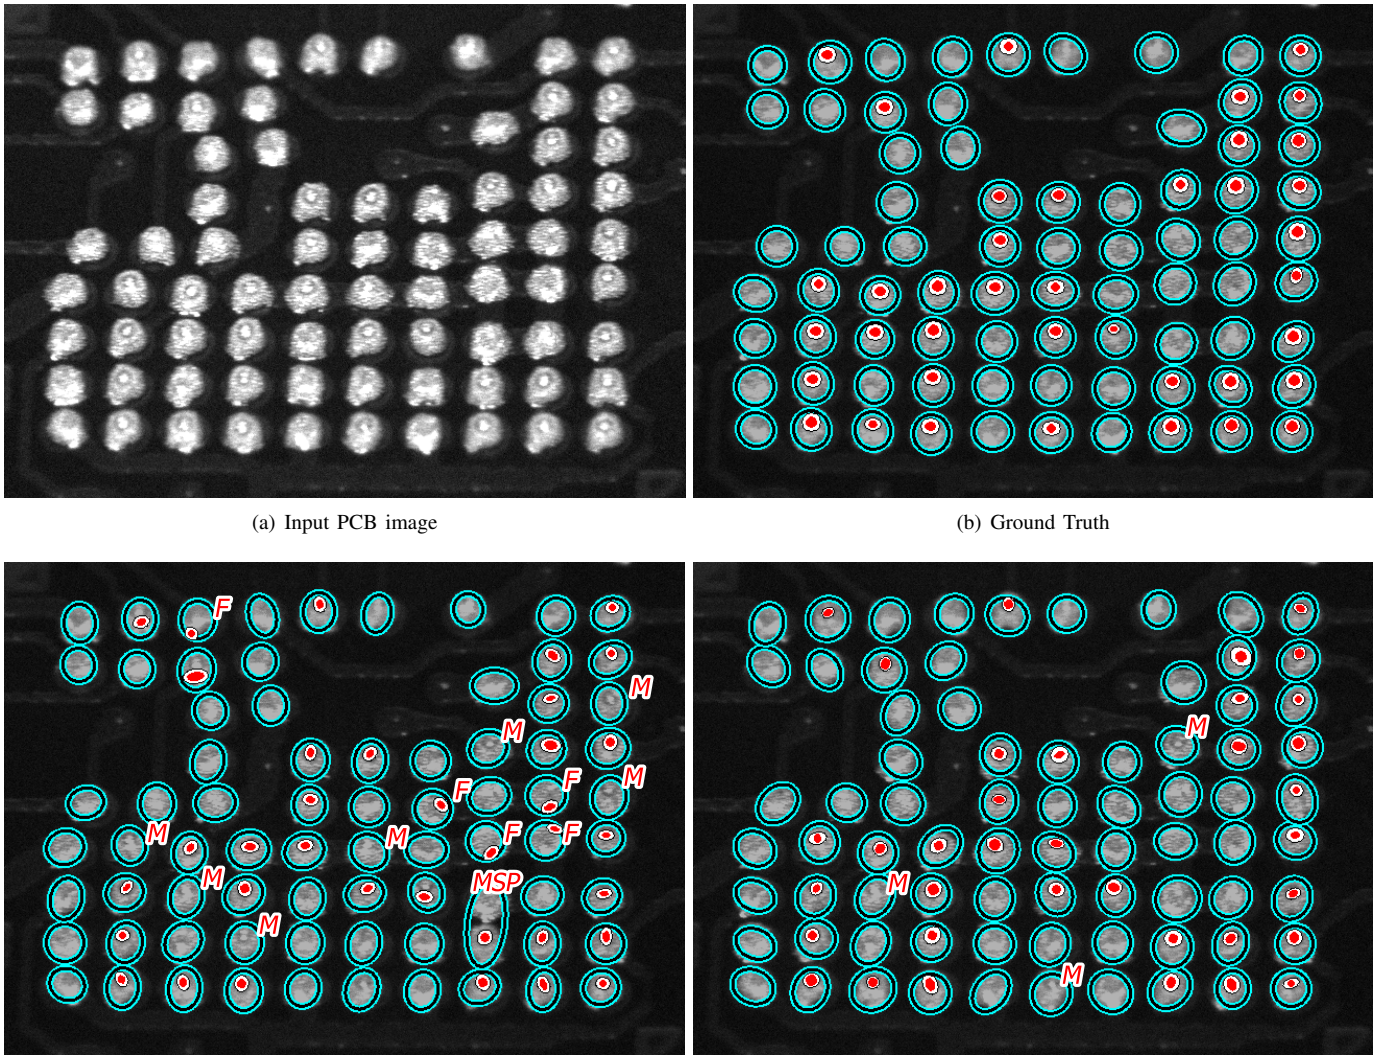

(a) Input PCB image

(b) Ground Truth

(c) Morph results: 2 merged SPs, 7 missing 10 false/falsely aligned SOs,

(d) HMPP results: 3 missing SOs

Fig. 12. Comparison of the detection results of the *Morph* and the proposed *HMPP* methods to the Ground Truth. *F* resp. *M* denote False resp. Missing scoops, *MSP* refers to an erroneously merged SP pair with the *Morph* technique.

Document version of the MTA SZTAKI Publication Repository, http://eprints.sztaki.hu/

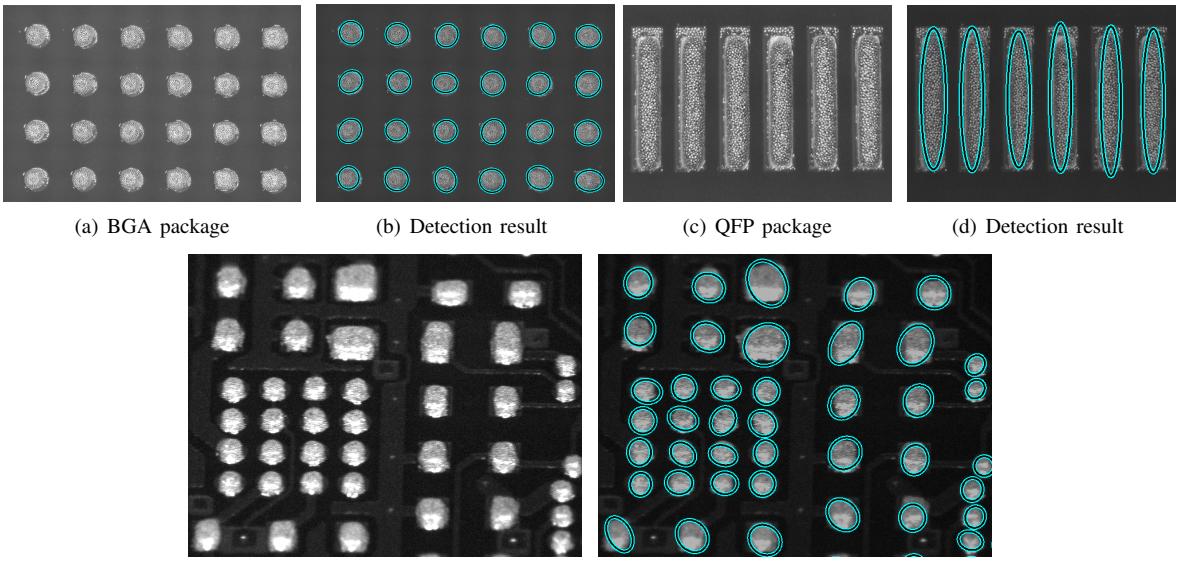

(a) BGA package (b) Detection result (c) QFP package (d) Detection result

(e) Nokia PCB data (f) Detection result

Fig. 13. SP detection results in different circuits

Document version of the MTA SZTAKI Publication Repository, http://eprints.sztaki.hu/