



The life and death of cosmic voids

P. M. Sutter, Pascal Elahi, Bridget Falck, Julian Onions, Nico Hamaus,
Alexander Knebe, Chaichalit Srisawat, Aurel Schneider

► To cite this version:

P. M. Sutter, Pascal Elahi, Bridget Falck, Julian Onions, Nico Hamaus, et al.. The life and death of cosmic voids. Monthly Notices of the Royal Astronomical Society, 2014, 445, pp.1235-1244. 10.1093/mnras/stu1845 . insu-03645234

HAL Id: insu-03645234

<https://insu.hal.science/insu-03645234>

Submitted on 25 Apr 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



The life and death of cosmic voids

P. M. Sutter,^{1,2,3★} Pascal Elahi,⁴ Bridget Falck,⁵ Julian Onions,⁶ Nico Hamaus,^{1,2}
Alexander Knebe,⁷ Chaichalit Srisawat⁸ and Aurel Schneider⁸

¹ Sorbonne Universités, UPMC Univ Paris 06, UMR7095, Institut d'Astrophysique de Paris, F-75014 Paris, France

² CNRS, UMR7095, Institut d'Astrophysique de Paris, F-75014 Paris, France

³ Center for Cosmology and AstroParticle Physics, Ohio State University, Columbus, OH 43210, USA

⁴ Sydney Institute for Astronomy, University of Sydney, Sydney, NSW 2016, Australia

⁵ Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth PO1 3FX, UK

⁶ School of Physics & Astronomy, University of Nottingham, Nottingham NG7 2RD, UK

⁷ Departamento de Física Teórica, Módulo C-15, Facultad de Ciencias, Universidad Autónoma de Madrid, E-28049 Cantoblanco, Madrid, Spain

⁸ Department of Physics & Astronomy, University of Sussex, Brighton BN1 9QH, UK

Accepted 2014 September 5. Received 2014 September 5; in original form 2014 March 28

ABSTRACT

We investigate the formation, growth, merger history, movement, and destruction of cosmic voids detected via the watershed transform code `VIDE` in a cosmological N -body dark matter Λ cold dark matter simulation. By adapting a method used to construct halo merger trees, we are able to trace individual voids back to their initial appearance and record the merging and evolution of their progenitors at high redshift. For the scales of void sizes captured in our simulation, we find that the void formation rate peaks at scale factor 0.3, which coincides with a growth in the void hierarchy and the emergence of dark energy. Voids of all sizes appear at all scale factors, though the median initial void size decreases with time. When voids become detectable they have nearly their present-day volumes. Almost all voids have relatively stable growth rates and suffer only infrequent minor mergers. Dissolution of a void via merging is very rare. Instead, most voids maintain their distinct identity as annexed subvoids of a larger parent. The smallest voids are collapsing at the present epoch, but void destruction ceases after scale factor 0.3. In addition, voids centres tend to move very little, less than 10^{-2} of their effective radii per $\ln a$, over their lifetimes. Overall, most voids exhibit little radical dynamical evolution; their quiet lives make them pristine probes of cosmological initial conditions and the imprint of dark energy.

Key words: cosmology: theory – large-scale structure of Universe.

1 INTRODUCTION

Since cosmic voids are, by definition, relatively empty of matter, they offer a unique and pristine laboratory for studying dark energy (Lavaux & Wandelt 2012; Sutter et al. 2012b), exotic fifth forces (Li, Zhao & Koyama 2012; Spolyar, Sahlén & Silk 2013), and the early universe (Goldberg & Vogeley 2004). They also offer a complementary probe of the growth of structure via their size and shape distributions (Biswas, Alizadeh & Wandelt 2010; Bos et al. 2012; Clampitt, Cai & Li 2013). Recently large catalogues of voids identified in galaxy redshift surveys (Pan et al. 2012; Sutter et al. 2012a; Nadathur & Hotchkiss 2014; Sutter et al. 2014c) have opened the way for statistical and systematic measurements of void properties (Ceccarelli et al. 2013; Sutter et al. 2014a), and their con-

nections to cosmological parameters (Melchior et al. 2014; Planck Collaboration XVI 2013).

However, given the promising utility of voids, we still lack a detailed understanding of their life cycles. For example, for a given void observed at low redshift, we do not know when it formed, where it formed, whether it grew to its present size via simple expansion or through mergers, nor whether it will continue expanding or eventually collapse. We also do not understand basic statistics about voids over cosmic time: their formation and merger rates, growth rates, and movement. Such understanding of the life cycles of voids will solidify current void-based cosmological analysis and enable future probes. Also, if we are to use voids as cosmological probes we must understand the impact of their dynamics on any primordial cosmological signal.

As identified theoretically by Sheth & van de Weygaert (2004) and discussed in the review of van de Weygaert & Platen (2011), void evolution appears intimately tied to its environment: smaller

★ E-mail: psutter2@illinois.edu

voids tend to appear inside larger overdense surroundings, while larger voids are truly anticorrelated with respect to the matter distribution. Thus, smaller voids tend to collapse over time, while larger voids continue to expand. The expansion of larger voids causes their interiors to appear as miniature open universes, with lower density walls, filaments, and haloes (van de Weygaert, Sheth & Platen 2004; Aragon-Calvo et al. 2010; Neyrinck et al. 2014) evolving in a self-similar pattern. The evolutionary and hierarchical behaviour has been modelled in the context of the adhesion approximation (Sahni, Sathyaprakash & Shandarin 1994), simulations (van de Weygaert & van Kampen 1993), and excursion set theory (Sheth & van de Weygaert 2004).

However, void abundances are still difficult to predict with excursion set formalisms alone (Jennings, Li & Hu 2013; Sutter et al. 2014b). While there have been several attempts to improve the initial theoretical result of Sheth & van de Weygaert (2004), such as by adjusting the void growth and destruction parameters (Furlanetto & Piran 2006; D’Aloisio & Furlanetto 2007; Paranjape, Lam & Sheth 2012) and rescaling void sizes (Jennings et al. 2013), there still remains very little correspondence to voids identified with watershed techniques in galaxy surveys (Sutter et al. 2014c). We may improve excursion set predictions by directly measuring the growth and destruction rate in cosmological simulations.

The shapes of voids offer a particularly interesting cosmological probe, whether by their distribution (e.g. Bos et al. 2012; Li et al. 2012) or via an application of the Alcock–Paczynski test (Alcock & Paczynski 1979; Ryden 1995; Lavaux & Wandelt 2012; Sutter et al. 2012b; Sutter, Pisani & Wandelt 2014d). However, these tests rely on the assumption that the void identified in a galaxy survey corresponds to a physical underdensity in the dark matter. While this is largely an issue of sparsity and galaxy bias (Sutter et al. 2014a), the watershed technique may spuriously merge voids even in the dark matter. These voids will erroneously appear as larger voids that are not completely empty and thus have suspect shapes. We can use a detailed merger history to identify such suspect voids.

Recently, Hamaus et al. (2014a) pointed out that for a given tracer population there exists a *compensation scale*, where the void-matter bias is identically zero. Below this scale, voids generally collapse due to their surrounding overdense walls, while above this scale voids tend to continue expanding (Ceccarelli et al. 2013). However, these results are based on studies of the velocity profiles and clustering statistics at fixed time. Only by tracing the evolution – and thereby studying the dynamics – of voids could one accurately examine the properties of voids in relation to such a compensation scale.

Finally, the growth and merger rates of voids are potential cosmological probes, analogous to the growth rate of cosmic structure. The nature of modified gravity and fifth forces can leave fingerprints on the evolution of the void population at high redshift, potentially constraining the properties of dark energy.

Unfortunately, to date this remains a largely unexplored topic. Most early studies of voids in simulations focused on visual identification and characterization (e.g. White et al. 1987). For example, the pioneering works of Dubinski et al. (1993), which discussed the process of void merging, and van de Weygaert & van Kampen (1993), which first noted the hierarchical nature of void buildup, were entirely based on visually examining thin slices of N -body simulations. More recent and more sophisticated analyses have focused on void interiors (e.g. Gottlober et al. 2003; Goldberg & Vogeley 2004; Aragon-Calvo & Szalay 2013; Neyrinck et al. 2014) or on statistics at a fixed time such as those discussed above.

In this work, we present a comprehensive study of the formation, subsequent evolution, and destruction of voids. We use techniques adapted from building halo merger trees (Srisawat et al. 2013) to follow individual voids across cosmic time. This approach allows us to measure their formation time, identify when mergers occur, track their movement and growth, and, when it does happen, record their time of collapse. We translate this information into rates and correlate these rates with void size, which can then inform theoretical and observational results.

In the following section, we review our simulation setup, void finding approach, merger identification technique, and some definitions to be used throughout the work. In Section 3, we focus on the formation time of voids, followed by a discussion in Section 4 of their growth and merger histories. In Sections 5 and 6, we present an analysis of void movement and destruction rate over cosmological time, respectively. Finally, we conclude in Section 7 with a brief discussion of implications for theoretical modelling of voids and directions for future work.

2 NUMERICAL APPROACH

2.1 Simulation

We study voids forming in a single cosmological dark matter simulation run using the GADGET-3 N -body code (Springel 2005) with initial conditions drawn from the WMAP-7 cosmology (Komatsu et al. 2011). Voids are identified in 62 snapshots from redshift 0 to ~ 30 . The snapshots are evenly spaced in $\ln a$, where a is the scale factor. The simulation contains 270^3 particles in a box of comoving length $62.5 h^{-1} \text{Mpc}$, giving a dark matter particle mass of $9.31 \times 10^8 h^{-1} M_\odot$. This combination of box size and number of particles gives a mean interparticle spacing of $\sim 0.25 h^{-1} \text{Mpc}$. Since we will only study voids of effective radius $1 h^{-1} \text{Mpc}$ (see below), this provides sufficient resolution for even the smallest voids and allows us to examine several thousand voids. Increasing resolution or box size would give us access to even more voids, but the analysis of Sutter et al. (2014a) shows that voids are self-similar up to a scale of $\sim 100 h^{-1} \text{Mpc}$ in a Λ cold dark matter universe: studying this distribution of voids gives us a fairly representative picture.

For more simulation details see Srisawat et al. (2013).

2.2 Void finding

We identify voids with a heavily modified and extended version of ZOBOV (Neyrinck 2008) called VIDE (Sutter et al., 2014e). VIDE creates a Voronoi tessellation of the tracer particle population and uses the watershed transform to group Voronoi cells into zones and subsequently voids (Platen, van de Weygaert & Jones 2007). By implicitly performing a Delauney triangulation (the dual of the Voronoi tessellation), VIDE assumes constant density across the volume of each Voronoi cell, which sets the smoothing scale for the continuous field necessary to perform the watershed transform. There is no additional smoothing.

The algorithm proceeds by first grouping adjacent Voronoi cells into zones, which are local basins. Next, the watershed transform merges zones into voids by examining the density barriers between them and joining them together to form ever-larger agglomerations. We impose a density-based threshold within VIDE where adjacent zones are only added to a void if the density of the wall between them is less than 0.2 times the mean particle density. Derived from the characteristic void non-linearity density level (Platen et al. 2007),




Figure 1. A cartoon of the assembly of the void hierarchy. The top panel shows ridgelines with line thickness proportional to density. The bottom panel shows the tree derived from such a collection of voids, with the tree level of each void indicated.

this prevents voids from expanding deeply into overdense structures and limits the depth of the void hierarchy (Neyrinck 2008). However, this does not place a restriction on the density of the initial zone, and in principle a void can have any mean density.

The watershed transform identifies catchment basins as the cores of voids and ridgelines, which separate the flow of water, as the boundaries of voids. In sum, we identify voids as depressions in the tracer density; voids are non-spherical aggregations of Voronoi cells that share a common basin and are bounded by a common set of higher density walls.

These operations allow the construction of a nested hierarchy of voids (Bos et al. 2012; Lavaux & Wandelt 2012): we identify the initial zones as the deepest voids, and as we progressively merge voids across ridgelines we can identify supervoids. There is no unique definition of a void hierarchy, and we take the semantics of Lavaux & Wandelt (2012): a parent void contains all the zones of a subvoid plus at least one more. All voids have only one parent but potentially many (or no) children, and the children of a parent occupy distinct subvolumes separated by low-lying ridgelines. Fig. 1 shows a cartoon of this void hierarchy construction. For visualiza-

tions of watershed voids, we refer the reader to Platen et al. (2007), Neyrinck (2008), Colberg et al. (2008), Bos et al. (2012), and Sutter et al. (2012a). Also, Aragon-Calvo et al. (2010) show the nested hierarchy of voids using different, but related, assembly criteria.

Our simulation gives a mean particle spacing $\bar{n}^{-1/3} \approx 0.25 h^{-1} \text{ Mpc}$, which sets a lower size limit of the detectability of voids due to shot noise. For this work, we will study all voids with effective radius $R_{\text{eff}} > 1 h^{-1} \text{ Mpc}$. We define the effective radius as

$$R_{\text{eff}} \equiv \left(\frac{3}{4\pi} V \right)^{1/3}, \quad (1)$$

where V is the total volume of the Voronoi cells that contribute to the void. Note that this cut is at a much larger radius than used in previous analyses with VIDE (e.g. Sutter et al. 2014c). While voids near the mean particle separation do not appear to be simple Poisson noise (Hamaus, Sutter & Wandelt 2014b), small-scale fluctuations in the density field can still give rise to occasional spurious features which are not filtered out, leading to an incomplete and inaccurate picture of the void population at these scales. However, at four times the mean particle separation, the abundance of voids appears complete (e.g. Sutter et al. 2014b), the void properties are convergent to higher resolutions (e.g. Sutter et al. 2014a), and the contamination by Poisson fluctuations is exponentially diminished (Neyrinck 2008; Sutter et al. 2013). Thus, for our study this criterion gives a rather robust picture of the void population. We do not impose any other cuts based on density contrast or minimum density – we wish to see if marginal voids have similar histories as deeper underdensities.

Additionally, for the analysis below we need to define a centre for each void. For this work we take the *macrocentre*, or volume-weighted centre of all the Voronoi cells in the void

$$\mathbf{x}_v = \frac{1}{\sum_i V_i} \sum_i \mathbf{x}_i V_i, \quad (2)$$

where \mathbf{x}_i and V_i are the positions and Voronoi volumes of each tracer i , respectively.

Fig. 2 shows the cumulative void number function for all voids in the final $a = 1.0$ snapshot, organized by level in the hierarchy. Note that these numbers functions do not turn over at small radii, as predicted by Sheth & van de Weygaert (2004), since we are plotting the cumulative, rather than differential, function, and we are only counting voids well above the resolution limit. At the topmost parent root level (Tree Level 0) there are only a few small field voids and the largest voids in the simulation. As we go deeper into the hierarchy, we see increasingly smaller voids. We construct the void tree such that each void has only a single parent (or no parents at all) and can potentially have many children. One void is a parent of another if it shares all zones of the child plus at least one more. Parents can then become children of even larger supervoids. Without any density thresholds, there will be a single void that encompasses the entire simulation volume. However, since we do apply a density threshold, we have multiple root voids. The relatively small simulation box prevents us from examining the very largest voids; however, the large voids that are discovered in this box are representative of the voids found in larger simulations and galaxy surveys.

2.3 Merger identification

To match voids from one snapshot to another, we use the tree building routine that is part of the publicly available VELOCIRAPTOR




Figure 2. Cumulative number functions for voids in different levels of the hierarchy. Tree level 0 (dark blue) are the topmost parent voids, and tree level 8 (red) are subvoids deepest in the hierarchy. Though higher level voids tend to be larger, they span a broad range of sizes. At the topmost level, there are only a few small ‘field’ (i.e. childless) voids and some voids that span nearly the entire simulation volume. We do not show these largest voids so that we may highlight the relative differences of the remaining hierarchy levels.

(aka STF) package.¹ The VELOCIRAPTOR tree builder code is a particle correlator: it takes two particle ID lists (named *A* and *B*) and for each object in list *B* identifies those objects in list *A* (i.e. in the previous snapshot) that have particles in common. This first step produces a graph mapping the connections between objects rather than an actual progenitor tree. To produce a tree, the algorithm calculates the merit of each connection

$$\mathcal{M}_{A_i B_j} = N_{A_i \cap B_j}^2 / (N_{A_i} N_{B_j}), \quad (3)$$

where $N_{A_i \cap B_j}$ is the number of shared particles between the object *i* in catalogue *A* and object *j* in catalogue *B*, and N_{A_i} and N_{B_j} are the number of particles in object *i* in catalogue *A* and *j* in *B*, respectively. Between two snapshots, the unique main progenitor is the object that maximizes the merit. This strategy has proven successful in building *halo* merger trees (see for instance Srisawat et al. 2013).

Note that while technically we are constructing bidirectional *graphs*, we will see that we are justified in calling these *trees* since the void evolution is surprisingly simple.

However, unlike haloes, which are defined by the particles they are composed of, voids are defined by the empty spaces between particles. Ideally, we would correlate *volumes* rather than *particles*. However, this is computationally expensive and fraught with difficulties: it would require modelling the Voronoi volume around each particle and making arbitrary decisions on when one particle’s volume is correlated with another. Instead of trying to determine the overlap in volume, we simply use the volume associated with each particle as determined by VIDE, v_i , to weight the merit function. Thus, the modified merit function becomes

$$\mathcal{M}_{A_i B_j} = \tilde{V}_{A_i \cap B_j}^2 / (V_{A_i} V_{B_j}), \quad (4)$$

where the total volume of a void is $V = \sum_l v_l$, and the shared volume is $\tilde{V}_{A_i \cap B_j} = \sum_l v_{l, A_i} v_{l, B_j}$. This sum is over all shared

particles, with each particle weighted by the volume associated with it in catalogues *A* and *B*. In order to qualify as a progenitor, the compared voids must share at least 10 particles, though changes to this value do not produce significantly different results.

We note that even a simple particle-based approach is justified as even our smallest voids contain hundreds of particles. While the cores of voids are empty, there are sufficient numbers of particle distributed throughout the remaining volume that particles can be used as a proxy for the void volume. The volume-weighting scheme used here simply reduces instances where a void in catalogue *A* shares particles with several voids in catalogue *B*. The volume-weighted merit function chooses the progenitor which minimises the volume fluctuations on a particle by particle basis and for the void as a whole.

In short, to construct a progenitor history for each void we find the void in the previous snapshot that shares the most number of particles, weighted by volume. This is called the *main progenitor*. We discuss the evolutionary chain of these main progenitors for the remainder of this work when we examine the movement and growth of a void.

3 VOID FORMATION

We begin with visual examination of voids as a function of scale factor, shown in Fig. 3. This figure shows slices of the dark matter particles that are identified as belonging to voids. Initially, there are only small, isolated voids just above the minimum size threshold. Since we do not apply a density criterion, these are shallow basins that will eventually empty out (as seen in van de Weygaert et al. 2004). These depressions in the initial density field then begin to expand, with small basins quickly merging with larger basins. As the walls and filaments begin to coalesce around $a = 0.2$ – 0.4 , a complex void hierarchy begins to form as subvoids group into larger parent points. At this epoch the density contrast in the large scale structure becomes high enough for our void finder to identify the multilevel basins as distinct voids. At late times, the larger voids simply grow and expand into their local environment, and since the watershed method includes all particles within the ridgeline as void members, we see very few gaps (e.g. the dense haloes) in the void particle distribution.

The shallow basins at early times could be considered as ‘proto-voids’. To separate these from mature voids one could define a density threshold, as is done for haloes (Sheth & van de Weygaert 2004). While linearly extrapolated initial conditions do not always clearly map to final void states (Sahni & Shandarin 1996), there is evidence that void *properties* exhibit only linear and quasi-linear behaviour. For example, the recent work of Hamaus et al. (2014b) showed that density and velocity profiles of voids, while each individually described by non-linear functions, can be related by linear perturbation theory very accurately. Also, Lavaux & Wandelt (2010) showed that it is possible to map the shapes of voids in Lagrangian initial conditions to the late-time Eulerian object. Thus, there is no clear distinction between early- and late-time voids: there is a one-to-one mapping from the initial to the final void state.

To clarify this mapping, Fig. 4 shows the formation scale factor a_f for each void identified at the present day. We define a_f as the scale factor at which the void reaches a fraction $f_v = V(a_f)/V_0$ of its current volume, V_0 . Since voids grow and shrink, defining a formation time is not trivial. We set a_f to the time at which the void’s volume drops below $f_v V_0$ for the *first* time as we go back along a void’s history. We indirectly investigate our definition of a formation time by examining several values of f_v . We use 10^{-4} ,

¹ <https://www.dropbox.com/sh/177zo6q3qk5pdzk/T0V0eseLZu>




Figure 3. A visual impression of the buildup of voids. Shown are thin slices of the particle distribution at various simulation snapshots. Only void member particles are shown. Slices are at scale factors 0.1 (left), 0.25 (middle), and 1.0 (right). Around $a = 0.23$, filaments and walls become dense enough and the underdensities clear enough to support the formation of larger voids. While the parent voids gently expand, the formation and merging of subvoids continues. Watershed void finders include as void members all particles within the highest density ridgeline; hence at late times almost all but the highest density particles are included in voids.




Figure 4. Distribution of scale factors of formation for voids that persist to $a = 1.0$. The peak around $a = 0.2$ – 0.4 corresponds to a significant increase in the depth of the void hierarchy.

8×10^{-3} , and 0.2, which appear chosen arbitrarily but pertain to specific increases in void size by corresponding to fractional radii of 0.01, 0.05, and 0.1, respectively. All three values give nearly identical results: some voids have persisted since the beginning of the simulation, and others have only appeared recently. As expected, with larger values of f_V the distribution skews to later times (note the much smaller value of the green line at $a = 0.0$), but this is surprisingly insignificant: once a void is detected, it has essentially its present-day volume. There is also an increase in the formation time just before $a = 1.0$. While this may be due to numerical effects, it is not significantly different than the $a_f > 0.4$ fluctuations.

There is a noticeable spike at $a \approx 0.3$. This coincides with the initial growth of the void hierarchy; prior to this time there are only isolated voids. Thus, this time indicates the appearance of significant structural hierarchy in the cosmic web. It also coincides with the appearance – though not domination – of dark energy: in these epochs, $\Omega_\Lambda \sim 0.1\Omega_M$. Thus, a large void population forms at these scale factors through a combination of the crystallization of the




Figure 5. Distribution of $a = 1$ void sizes as a function of their formation scale factor a_f . The black line is the median in bins of width $\Delta a_f = 0.05$, the dark bands are the inner 68 per cent of the binned distribution, and light grey bands are the bin extrema.

cosmic web, allowing basins to be identified, and the emergence of dark energy, which shuts off significant continued void production. Since all chosen values of f_V give nearly the same results, voids reach their present-day volume quickly – essentially when they are first able to be identified as voids – and do not grow much. We will return to this subject later.

In Fig. 5, we examine the distribution of void sizes as they appear in the simulation. Throughout most of cosmic history, the median void formation size is centred on $2 h^{-1}$ Mpc, and a small population of medium-scale 10 – $15 h^{-1}$ Mpc voids continually appears. The median void formation drops to nearly $1 h^{-1}$ Mpc by the present day, since newer voids can only occupy smaller niches in the cosmic web adjacent to larger, expanding neighbour voids. The very largest voids appear at any times, as the walls between mid-scale voids empty out and the voids are joined into larger supervoids. The fact that there is no noticeable peak formation time for voids of a specific size but there is a peak in the formation time of voids indicates that the level of these voids in the hierarchy changes with time. Although we are showing the present-day $a = 1$ void sizes, the fact that voids do not grow much over time indicates that this is also essentially their size at formation.

4 VOID MERGERS AND GROWTH

We use Fig. 6 to show an example of a void merger history. In this figure, we show slices of the dark matter density centred on a representative void at $a = 1.0$. We then show the same slice at the same position at two other scale factors, 0.8 and 0.4. In these additional slices, we show the progenitors of the present-day void. While the voids in general have complex shapes due to the nature of the watershed (see for example, Sutter et al. 2012a), we represent them here as simple circles with radii equal to the effective radius R_{eff} ; our purpose with this illustration is to simply provide a rough guide to the eye of void location and size relative to the cosmic web, not to examine the detailed impact on void shapes.

We choose this void to highlight two distinct processes of void evolution. The first, pure merging, occurs when high-density barriers between two voids completely dissolve due to outflows. When the barrier becomes too low ($< 0.2\bar{\rho}$), the separate voids become indistinguishable from each other, forming a single larger void. However, as we will discuss below, this is a very rare process. Instead, what more frequently occurs is *annexation* of subvoids as a larger parent void assembles, generating a hierarchy. In this picture, a subvoid undergoing annexation retains its detectability and individual identity, and comprises only a portion of the volume of a larger parent (Fig. 1). In contrast, a *merging* subvoid loses its identity and is no longer separately detectable.

For each progenitor tree leading up to each present-day void, we can track the total number of progenitors; in other words, the width of the merger tree. In Fig. 7, we plot the average number of progenitors as a function of the scale factor for all present-day voids,

$$\langle N_{\text{prog}} \rangle(a) = \frac{1}{N_{\text{voids},0}} \sum_i N_{\text{prog},i}(a). \quad (5)$$

Here, the sum is over all progenitors at a time a . The fact that the mean number of progenitors is barely greater than one indicates that almost all voids follow only a single line of descent and experience very few mergers. The number of progenitors begins to decrease at $a = 0.4$, which coincides with the end of the peak formation time (Fig. 4) and the relative lack of new voids after that time.




Figure 6. Evolution of progenitor voids. We show thin slices through the dark matter density with voids superimposed on top. Voids are represented as circles with radii equal to R_{eff} . Slices are arranged from early (left) to late (right) times and trace the evolution of a single void to highlight different void merger histories. We only show progenitors of the final $a = 1.0$ void. To increase the density contrast, each panel is scaled such that black and white are the minimum and maximum density of the cells shown, respectively. The density is constructed using cloud-in-cell weighting and the shading is scaled according to $\log(1 + \delta)$. Projection effects lead to void centres occasionally appearing to lie on top of filaments. Several processes are highlighted by this evolution, including void merging as barriers dissolve, and the formation of a void hierarchy as the larger parent void annexes smaller, but distinct, subvolumes.




Figure 7. Average number of progenitors as a function of scale factor for voids at $a = 1.0$. We calculate this quantity by counting the total number of progenitors for all present-day voids and dividing by the number of present-day voids. Hence, this quantity becomes less than one as the lines of descent for individual voids end.

In Fig. 8, we show the growth rate history for every void in the simulation as a function of scale factor a ,

$$\frac{d \ln R_{\text{eff}}}{d \ln a} = \ln(R_{i+1}/R_i) / \ln(a_{i+1}/a_i), \quad (6)$$

where the change is calculated between the i th + 1 and i th snapshots. We plot a line for each individual void. As expected, the voids with the highest growth rates are the largest; these are the supervoids that form from the rapid merger of subvoids as the basin empties out of substructure. Since they form in deeply underdense environments with little substantial structure surrounding them, they act as




Figure 8. Growth rate as a function of scale factor for each void surviving at $a = 1.0$. Individual lines are coloured by void size, from the smallest (blue) to largest (red). The radii listed are the final $a = 1.0$ size.

miniature universes with $\Omega_{\text{tot}} < 1$ (Goldberg & Vogeley 2004). For these larger voids, there are a few steep changes as they merge with a smaller void or fragment into smaller progenitors. Even though the merger or fragmentation ratio is small in terms of volume, it can impact the instantaneous growth rate from one snapshot to another.

The small- and medium-scale voids show remarkably steady growth histories, with very few strong deviations. Even though some of these voids do merge, they tend to just absorb their sub-

voids, so the overall volume gained is small. Interestingly, there is a population of collapsing voids: these are the voids located in overall overdense regions. This is the ‘void-in-cloud’ phenomenon of Sheth & van de Weygaert (2004). Even though there are a few small voids with discontinuous merger histories, almost all the small voids are either gently expanding or contracting.

We break down the growth rates into secular and merging components, as we show in Fig. 9. We define the secular growth rate as the growth rate of voids which did not experience a merger in that timestep. If instead that void did merge with another, its growth from snapshot to snapshot is calculated in the average merger growth rate. Here, we plot the mean growth rate over all voids as a function of scale factor. We also separate voids into their level in the hierarchy so that we may examine the nature of larger parent voids and their subvoids separately.

First, we note that the merger growth rate far outweighs the secular growth rate by an order of magnitude; voids gain volume typically not by growth of the underlying volume but by merging (when it does occur) with adjacent voids. The fractional merger growth rate is much larger for subvoids deep in the hierarchy than it is for voids higher in the tree. Thus, even though a small fraction of larger voids experience occasional large jumps in their volume, averaged over the entire cohort of voids in that tree level it is entirely insubstantial. We see for voids deep in the hierarchy (that is, subvoids) a peak in the merger growth rate at scale factor 0.3, which is the epoch with the highest formation rate where voids experience a rapid restructuring as the hierarchy forms.

Comparatively, the secular growth rate is very small and contributes little to the overall growth of voids. Here, we see the opposite trend as for the merger-based growth rate: the voids highest in the tree hierarchy have the highest rates, since they are not surrounded by overdense shells that would restrict their growth. Again we see a collapsing void population. Indeed, they have been collapsing since $a = 0.3$, when the top-level voids first formed.

Finally, Fig. 10 shows the instantaneous growth rate at $a = 1.0$. This is an interesting statistic since it allows us to separate collapsing

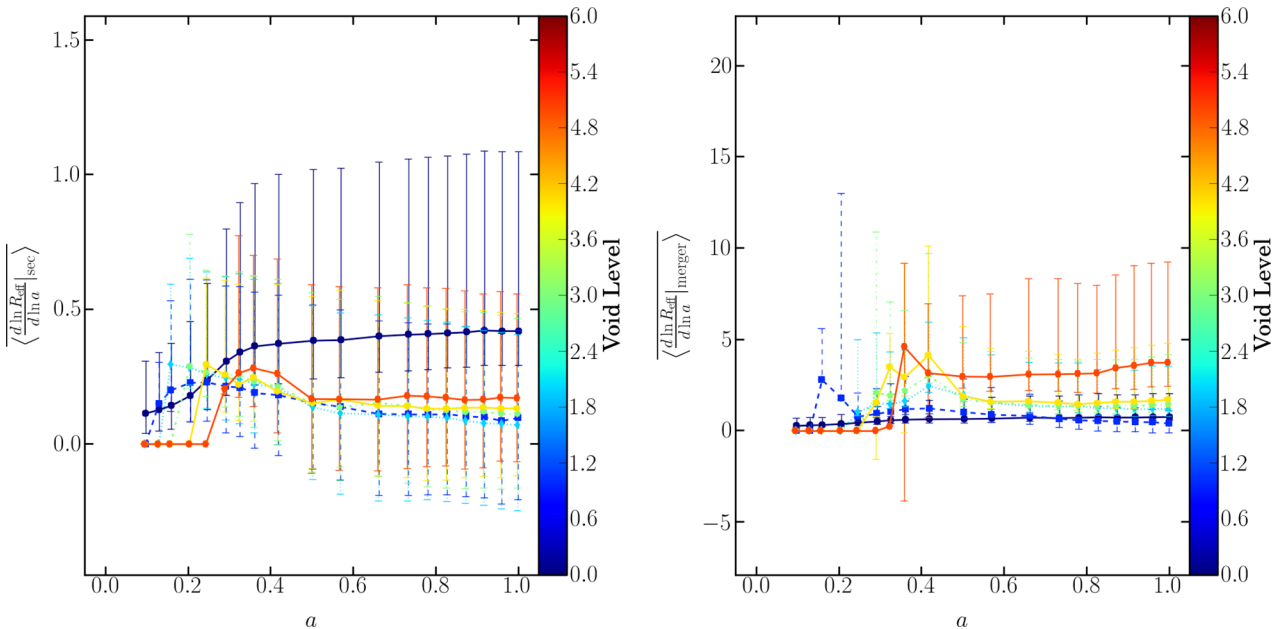


Figure 9. Breakdown of void growth rate into secular (left-hand panel) and merging (right-hand panel) components. Shown is the mean growth rate (solid lines) and the maximum spread (dashed error bars) for all voids as a function of scale factor. Voids are grouped by their position in the void hierarchy, from the topmost parent voids (blue) to the deepest children void (red).




Figure 10. Instantaneous secular (i.e. not merging) growth rate at $a = 1.0$ as a function of void effective radius. The solid black line is the median growth rate in bins of $1 h^{-1}$ Mpc, while the dark and light bands are the 68 and 95 per cent binned percentiles in the distributions.

from expanding voids based purely on their recent history. Similar analyses have been done using the void–matter cross-correlation (Hamaus et al. 2014a) and the identification of universal density and velocity profiles (Hamaus et al. 2014b; Sutter et al. 2014b). While we do not have sufficient volume to reliably apply those measures here, we do see a transition to overall-collapsing voids around $1\text{--}2 h^{-1}$ Mpc, although there are subsets of collapsing voids at all but the largest scales. A cleaner separation is based on the position in the void hierarchy, as discussed above. The clusterings statistics of Hamaus et al. (2014a) estimate a compensation scale roughly around $1 h^{-1}$ Mpc, but this is very uncertain due to the small box size. Despite this, we see very good correspondence between that result and the point at which the voids, on average, are collapsing.

5 VOID MOVEMENT

Despite the occasional violent merger, all voids experience very little movement of their macrocentres (equation 2). Fig. 11 shows the mean void macrocentre velocity, which we define as $\langle d|X|/d \ln a \rangle$, and which we express as fractions of the void effective radius at the current epoch. The average is taken over the entire void lifetime tracing from its current state along the branch of its main progenitors.

The distribution of mean macrocentre velocities has two distinct peaks. One, at 10^{-3} , represents the movement of the majority of voids, and is remarkably low. This mean velocity gives rise to an average displacement of only a tiny fraction of the void radius. This is not surprising: once a deep underdensity forms from the initial conditions, it is unlikely to move as the voids expand into the surrounding cosmic web.

The second peak, at 1×10^{-2} , is the small population of voids that experience violent mergers. Most of these voids are small subvoids residing deep in the void hierarchy. As before, we see that this is only a small fraction, roughly five per cent, of all voids. Even in these cases, the mean displacement is very small, indicating that even when voids experience mergers they are relatively gentle, since the macrocentre remains relatively stable.




Figure 11. Distribution of mean velocities of the void macrocentre over the void lifetimes.

6 VOID DESTRUCTION

We show the destruction rate, or the fraction of the existing void population lost in each snapshot in Fig. 12. While initially high prior to $a = 0.3$, afterwards the destruction rate drops to essentially zero. The initial relatively high destruction rate is not surprising, since at these early times the basins are just beginning to become deep enough to be identified as voids, and there is significant noise in the classification of these objects. However, even at these low-scale factors no more than ~ 4 per cent of voids are lost in every snapshot. But at the same time the formation rate spikes, $a = 0.3$, the destruction rate plummets. After this epoch, voids that have already




Figure 12. Fraction of the existing void population destroyed as a function of scale factor.

existed or will eventually form never cease to exist. Thus, even voids that are collapsing are not squeezed entirely; instead they merely become subvoids of larger parent voids. In fact, the squeezing may help bolster their ability to be detected: as the overdense shells around them grow higher, the density contrasts increase, allowing the void finder to continually detect them.

7 CONCLUSIONS

We have performed a comprehensive analysis of the life cycle – covering formation, mergers, growth, movement, and destruction – of cosmic voids. We have adapted merger tree codes originally designed to track the evolution of haloes to account for the large spatial extents of voids. By applying this technique to a high-resolution N -body simulation, we have gained a clear picture of voids as dynamic objects in the cosmic web. Through the use of a watershed void finder, we are able to classify voids according to their position in a hierarchy and use that to identify key epochs and scales in their evolution.

The past life of a cosmic void depends intimately on its place in the void hierarchy. Voids near the top of the hierarchy primarily form at a scale factor of 0.3, when the density contrasts in the cosmic web become high enough to support their identification and the introduction of dark energy shuts off continued structure formation. These higher level voids suffer only minor mergers and tend to maintain consistent growth rates over cosmic time. In contrast, voids lying deep in the hierarchy continue to form and have a somewhat more violent life due to the lower density nature of their surroundings, but even most of these voids have only a single line of descent. The location of a void in the hierarchy is more important than its size: two voids of equal volume can have radically different merger histories depending on their amount of substructure.

Voids typically grow at slow rates. However, there is a population of small collapsing voids. These voids tend to live in overdense environments near filaments and walls, as initially pointed out by van de Weygaert et al. (2004). Their overdense surroundings slowly squeeze them as adjacent larger voids expand. This picture is consistent with the theory developed by Sheth & van de Weygaert (2004), the velocity inflow–outflow analysis of Ceccarelli et al. (2013), the clustering study of Hamaus et al. (2014a), and the density profile studies of Hamaus et al. (2014b) and Sutter et al. (2014b). Despite being slowly crushed, after $a = 0.3$, these voids never get completely destroyed. Instead, they continue to survive as identifiable voids to the present day. Thus, the void destruction rate does not play a significant role in the late-time evolution of voids, and can be ignored in theoretical treatments. Additionally, as pointed out by Russell (2013), the collapsing process is completely negligible for all but the smallest voids, although this result is in conflict with analyses based on the adhesion approximation (Sahni et al. 1994).

Finally, voids do not move much throughout their lifetimes. Only small voids in the frothy depths of the hierarchy that undergo several mergers appear to have perturbed macrocentres. Even for these most active of voids, they typically only move a few per cent of their effective radii.

The combination of small box volume and high-resolution limits our study to relatively small (~ 1 – $15 h^{-1}$ Mpc) voids, while voids in larger simulations and galaxy surveys are typically much larger. However, recently Sutter et al. (2014b) were able to show that many void properties scale as a function of sampling density and galaxy bias. Thus, properties and characteristics of voids studied in one population of tracers can, in principle, be immediately translated to voids in another population. Thus, the conclusions that we reach

in this work are generally applicable to voids discovered in other simulations and galaxies.

We have examined the properties of voids defined using a watershed technique. There are, of course, other plausible definitions of voids (see for example, the comparison work of Colberg et al. 2008). These different algorithms might give different pictures of void histories, especially formation times, since they usually impose density thresholds. Additionally, there are other approaches to defining merger trees. We have noticed that volume correlations based on particles can give some non-intuitive results: voids that appear to occupy similar positions (based on their macrocentres and effective radii) may not necessarily share any particles. The relationships between particle correlation and macrocentre definition should be investigated further. However, we have applied other merger tree algorithms, such as *MergerTree* and *JMerge* (both described in Srisawat et al. 2013), and found qualitatively similar results.

Overall, voids live far quieter lives than their overdense counterparts, the haloes. Whereas up to 20 per cent of haloes have suffered a recent major merger, voids experience essentially *no* major mergers throughout their lifetime. Likewise, while subhaloes can be stripped of their mass as they pass through a larger parent halo, subvoids continue to be identifiable even when a supervoid forms around them. The implication is that voids are a much more pure cosmological probe; the fundamental cosmological signal imprinted from initial conditions and modified by dark energy is not corrupted by significant dynamics. Thus lower redshift cosmological probes, such as the Alcock–Paczynski test and void–galaxy cross-correlations, will not be affected by recent spurious mergers in the void population.

ACKNOWLEDGEMENTS

PMS acknowledges support from NSF Grant NSF AST 09-08693 ARRA, and would like to thank Mark Neyrinck, Alice Pisani, Ravi Sheth, Ben Wandelt, and David Weinberg for useful comments and discussions.

PJE would like to thank Krzysztof Bolejko for many useful discussions over tea. PJE is supported by the SSIMPL programme and the Sydney Institute for Astronomy (SIfA), DP130100117.

BF is supported by the UK Science and Technology Facilities Council grants ST/K00090/1 and ST/L005573/1.

AK is supported by the *Ministerio de Economía y Competitividad* (MINECO) in Spain through grant AYA2012-31101 as well as the Consolider-Ingenio 2010 Programme of the *Spanish Ministerio de Ciencia e Innovación* (MICINN) under grant MultiDark CSD2009-00064. He also acknowledges support from the *Australian Research Council* (ARC) grants DP130100117 and DP140100198. He further thanks Momus for the myth of the bishonen.

All the authors are thankful for the generosity and hospitality of the organizers of the Sussing Merger Trees Workshop,² held in 2013 July, where this project started, and especially Frazer Pearce and Peter Thomas, who provided valuable comments. The Sussing Merger Trees Workshop was supported by the European Commission’s Framework Programme 7 through the Marie Curie Initial Training Network CosmoComp (PITN-GA-2009-238356). This also provided fellowship support for AS.

The authors contributed in the following ways to this paper: AK, CS, and AS were members of the SOC that organized the workshop from which this project originated. They helped design

² <http://popia.ft.uam.es/SussingMergerTrees>

the comparison, plan, and organize the data. The analysis presented here was performed by PE, PMS, BF, JO, and NH and the paper was written by PMS. The other authors contributed towards the content of the paper and helped to proofread it.

REFERENCES

- Alcock C., Paczynski B., 1979, *Nature*, 281, 358
- Aragon-Calvo M. A., Szalay A. S., 2013, *MNRAS*, 428, 3409
- Aragon-Calvo M. A., van de Weygaert R., Araya-Melo P. A., Platen E., Szalay A. S., 2010, *MNRAS*, 404, L89
- Biswas R., Alizadeh E., Wandelt B., 2010, *Phys. Rev. D*, 82
- Bos E. G. P., van de Weygaert R., Dolag K., Pettorino V., 2012, *MNRAS*, 426, 440
- Ceccarelli L., Paz D., Lares M., Padilla N., Lambas D. G., 2013, *MNRAS*, 434, 1435
- Clampitt J., Cai Y.-C., Li B., 2013, *MNRAS*, 431, 749
- Colberg J. M. et al., 2008, *MNRAS*, 387, 933
- D’Aloisio A., Furlanetto S. R., 2007, *MNRAS*, 382, 860
- Dubinski J., da Costa L. N., Goldwirth D. S., Lecar M., Piran T., 1993, *ApJ*, 410, 458
- Furlanetto S. R., Piran T., 2006, *MNRAS*, 366, 467
- Goldberg D. M., Vogeley M. S., 2004, *ApJ*, 605, 1
- Gottlober S., Lokas E. L., Klypin A., Hoffman Y., 2003, *MNRAS*, 344, 715
- Hamaus N., Wandelt B. D., Sutter P. M., Lavaux G., Warren M. S., 2014a, *Phys. Rev. Lett.*, 112, 041304
- Hamaus N., Sutter P. M., Wandelt B. D., 2014b, *Phys. Rev. Lett.*, 112, 251302
- Jennings E., Li Y., Hu W., 2013, *MNRAS*, 434, 2167
- Komatsu E. et al., 2011, *ApJS*, 192, 18
- Lavaux G., Wandelt B. D., 2010, *MNRAS*, 403, 1392
- Lavaux G., Wandelt B. D., 2012, *ApJ*, 754, 109
- Li B., Zhao G.-B., Koyama K., 2012, *MNRAS*, 421, 3481
- Melchior P., Sutter P. M., Sheldon E., Krause E., Wandelt B. D., 2014, *MNRAS*, 440, 2922
- Nadathur S., Hotchkiss S., 2014, *MNRAS*, 440, 1248
- Neyrinck M. C., 2008, *MNRAS*, 386, 2101
- Neyrinck M. C., Aragón-Calvo M. A., Jeong D., Wang X., 2014, *MNRAS*, 441, 646
- Pan D. C., Vogeley M. S., Hoyle F., Choi Y.-Y., Park C., 2012, *MNRAS*, 421, 926
- Paranjape A., Lam T. Y., Sheth R. K., 2012, *MNRAS*, 420, 1648
- Planck Collaboration XVI, 2013, preprint ([arXiv:1303.5079](https://arxiv.org/abs/1303.5079))
- Platen E., van de Weygaert R., Jones B. J. T., 2007, *MNRAS*, 380, 551
- Russell E., 2013, *MNRAS*, 436, 3525
- Ryden B. S., 1995, *ApJ*, 452, 25
- Sahni V., Shandarin S., 1996, *MNRAS*, 282, 641
- Sahni V., Sathyaprakash B. S., Shandarin S. F., 1994, *ApJ*, 431, 20
- Sheth R. K., van de Weygaert R., 2004, *MNRAS*, 350, 517
- Spolyar D., Sahlén M., Silk J., 2013, *Phys. Rev. Lett.*, 111, 241103
- Springel V., 2005, *MNRAS*, 364, 1105
- Srisawat C. et al., 2013, *MNRAS*, 436, 150
- Sutter P. M., Lavaux G., Wandelt B. D., Weinberg D. H., 2012a, *ApJ*, 761, 44
- Sutter P. M., Lavaux G., Wandelt B. D., Weinberg D. H., 2012b, *ApJ*, 761, 187
- Sutter P. M., Lavaux G., Wandelt B. D., Weinberg D. H., 2013, preprint ([arXiv:1310.5067](https://arxiv.org/abs/1310.5067))
- Sutter P. M., Lavaux G., Wandelt B. D., Weinberg D. H., Warren M. S., 2014a, *MNRAS*, 438, 3177
- Sutter P. M., Lavaux G., Hamaus N., Wandelt B. D., Weinberg D. H., Warren M. S., 2014b, *MNRAS*, 442, 462
- Sutter P. M., Lavaux G., Wandelt B. D., Weinberg D. H., Warren M. S., Pisani A., 2014c, *MNRAS*, 442, 3127
- Sutter P. M., Pisani A., Wandelt B. D., 2014d, *MNRAS*, 443, 2983
- Sutter P. M. et al., 2014e, preprint ([arXiv:1406.1191](https://arxiv.org/abs/1406.1191))
- van de Weygaert R., Platen E., 2011, *Int. J. Modern Phys. Conf. Ser.*, 1, 41
- van de Weygaert R., van Kampen E., 1993, *MNRAS*, 263, 481
- van de Weygaert R., Sheth R., Platen E., 2004, in Diaferio A., ed., *IAU Colloq. 195: Outskirts of Galaxy Clusters: Intense Life in the Suburbs*. Cambridge Univ. Press, Cambridge, p. 58
- White S. D. M., Frenk C. S., Davis M., Efstathiou G., 1987, *ApJ*, 313, 505

This paper has been typeset from a \LaTeX file prepared by the author.