Coleman, G., Davín, A. A., Mahendrarajah, T., Szantho, L., Spang, A., Hugenholtz, P., Szöllősi, G. J., & Williams, T. (2021). A rooted phylogeny resolves early bacterial evolution. *Science*, 372(6542), [eabe0511]. <https://doi.org/10.1126/science.abe0511>

Peer reviewed version

Link to published version (if available):
[10.1126/science.abe0511](https://doi.org/10.1126/science.abe0511)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via American Association for the Advancement of Science at <https://doi.org/10.1126/science.abe0511> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

A rooted phylogeny resolves early bacterial evolution

Gareth A. Coleman^{1*}, Adrián A. Davín^{2*}, Tara A. Mahendrarajah³, Lénárd L. Szánthó^{4,6}, Anja Spang^{3,5}, Philip Hugenholtz²⁺, Gergely J. Szöllősi^{4,6,7+}, Tom A. Williams¹⁺

1. School of Biological Sciences, University of Bristol, Bristol BS8 1TH, UK.
2. Australian Centre for Ecogenomics, School of Chemistry and Molecular Biosciences, The University of Queensland, Queensland, Australia.
3. NIOZ, Royal Netherlands Institute for Sea Research, Department of Marine Microbiology and Biogeochemistry; and Utrecht University, AB Den Burg, The Netherlands
4. Dept. of Biological Physics, Eötvös Loránd University, 1117 Budapest, Hungary
5. Department of Cell- and Molecular Biology, Science for Life Laboratory, Uppsala University, SE-75123, Uppsala, Sweden
6. MTA-ELTE “Lendület” Evolutionary Genomics Research Group, 1117 Budapest, Hungary;
7. Institute of Evolution, Centre for Ecological Research, 1121 Budapest, Hungary

*Contributed equally to this paper.

+ Co-corresponding authors; contributed equally to this paper; correspondence to p.hugenholtz@uq.edu.au, ssolo@elte.hu, tom.a.williams@bristol.ac.uk

A rooted bacterial tree is necessary to understand early evolution, but the position of the root is contested. Here, we model the evolution of 11,272 protein families to identify the root, extent of horizontal gene transfer (HGT), and the nature of the last bacterial common ancestor (LBCA). Our analyses root the tree between the major clades Terrabacteria and Gracilicutes and suggest that LBCA was a free-living flagellated, rod-shaped double-membraned organism. Contrary to recent proposals, our analyses reject a basal placement of the Candidate Phyla Radiation, which instead branches sister to Chloroflexota within Terrabacteria. While most (92%) gene families have evidence of HGT, overall two-thirds of gene transmissions have been vertical, suggesting that a rooted tree provides a meaningful frame of reference for interpreting bacterial evolution.

A species tree captures the relationships among organisms but requires a root to provide the direction of evolution. Rooting deep radiations (1) is among the greatest challenges in phylogenetics, and there is no consensus on the root of the bacterial tree. On the basis of evidence (2–5) that the root of the tree of life lies between Bacteria and Archaea, early analyses with an archaeal outgroup placed the bacterial root near Aquificales/Thermotogales (6, 7) or Planctomycetes (8). Alternative approaches, including analyses of gene flows and the evolution of multimeric protein complexes as well as other complex characters (9), have instead suggested roots within the monoderm (single-membrane) Bacteria (10), or between Chloroflexi and all other cellular life (9). The development of techniques for sequencing microbes directly from environmental samples, without the need for laboratory cultivation, has greatly expanded the genomic representation of natural prokaryotic diversity (11–14). Recent phylogenomic analyses of expanded sets of diverse bacteria have placed the root between

one of the recently identified groups, the Candidate Phyla Radiation (CPR; also known as Patescibacteria (15, 16)) and all other Bacteria (11, 16, 17). The CPR is characterized by small cells and genomes that appear to have predominantly symbiotic or parasitic lifestyles, but much remains to be learned about their ecology and physiology (15, 17–19). If correct, the early divergence of the CPR has important implications for our understanding of the earliest period of cellular evolution. Along with evidence that the root of the archaeal domain lies between the reduced and predominantly host-associated DPANN superphylum and all other Archaea (1, 20), the CPR root implies that streamlined, metabolically minimalist prokaryotes have co-existed with the more familiar, self-sufficient lineages throughout the history of cellular life (19, 21).

Improved taxon sampling can help to resolve evolutionary relationships (22, 23), and the quantity and diversity of genome sequence data now available presents an opportunity to address long-standing questions about the origins and diversification of Bacteria. However, deep phylogenetic divergences are difficult to resolve, both because the phylogenetic signal for such relationships is overwritten by new changes through time, and also because the process of sequence evolution is more complex than the best-fitting models currently available. In particular, variation in nucleotide or amino acid composition across the sites of the alignment and the branches of the tree can induce long branch attraction (LBA) artifacts in which deep-branching, fast-evolving, poorly-sampled or compositionally biased lineages group together irrespective of their evolutionary history (24). These issues are widely appreciated (11) but are challenging to address adequately, particularly when sequences from thousands of taxa (11, 13, 14, 16, 17) are used to estimate trees of global prokaryotic diversity, which precludes the use of the best-fitting phylogenetic methods available.

Archaeal outgroup rooting does not unambiguously establish the root of the bacterial tree

The standard approach to rooting is to include an outgroup in the analysis, and all published phylogenies in which CPR form a sister lineage to the rest of the Bacteria (11, 16, 17) have made use of an archaeal outgroup. Outgroup rooting on the bacterial tree, however, has three serious limitations. First, interpretation of the results requires the assumption that the root of the tree of life lies between Bacteria and Archaea. While this is certainly the consensus view, the available evidence is limited and difficult to interpret (2–5, 25), and alternative hypotheses in which the universal root is placed within Bacteria have been proposed on the basis of indels (26, 27) or the analysis of slow-evolving characters (9). Second, the long branch leading to the archaeal outgroup has the potential to distort within-Bacteria relationships because of LBA. Third, joint analyses of Archaea and Bacteria use a smaller number of genes that are widely conserved and have evolved vertically since the divergence of the two lineages, and sequence alignment is more difficult due to the low sequence identity between homologues of the two domains.

We evaluated the performance of outgroup rooting on the bacterial tree using 143 Archaea and a shared subset of 29 phylogenetic markers (table S1). Using this archaeal outgroup, the maximum likelihood (ML) phylogeny under the best-fitting model (LG+C60+R8+F, which accounts for site-heterogeneity in the substitution process) placed the bacterial root between a clade comprising Cyanobacteria, Margulisbacteria, CPR, Chloroflexota and Dormibacterota on one side of the root, and all other taxa on the other (fig. S1). However, bootstrap support for this root, and indeed many other deep branches in both the bacterial and archaeal subtrees

was low (50-80%). We therefore used approximately-unbiased (AU) tests (28) to determine whether a range of published alternative rooting hypotheses (table S2) could be rejected, given the model and data. The AU test asks whether the optimal trees that are consistent with these other hypotheses have a significantly worse likelihood score than the maximum likelihood tree. In this case, the likelihoods of all tested trees were statistically indistinguishable ($AU > 0.05$, table S2), indicating that outgroup rooting cannot resolve the bacterial root on this alignment.

An alternative to outgroup rooting for deep microbial phylogeny

Given the limitations of using a remote archaeal outgroup to establish the root of the bacterial tree, we explored outgroup-free rooting using gene tree-species tree reconciliation (1, 29–31). We recently applied this approach to root the archaeal tree (1), and similar approaches have been used to investigate the root of eukaryotes (32, 33) and to map and characterize whole genome duplications in plants (34). Gene tree-species tree reconciliation methods work by adding a layer to the standard framework for inferring trees from molecular data. This additional step models the way in which gene trees can differ from each other and the overarching rooted species tree. Substitution models (such as LG(35)) describe how the constituent sequences of a gene family evolve along a gene tree via a series of amino acid substitutions, and allow us to infer the most likely gene tree. Reconciliation models describe how a gene tree evolves along the rooted species tree, beginning with gene birth (origination) and followed by a combination of vertical descent and events such as gene duplications, transfers, and losses (this series of events is a reconciliation). Combining the substitution-based modelling of sequences along the gene tree with the reconciliation-based modelling of gene trees along a rooted species tree allows us to infer the most likely rooted species tree from the constituent gene families. In other words, reconciliation methods aggregate phylogenetic signal across gene families and, since the likelihood of reconciliations depends on the position of the root, can be used to test the support for competing root positions (1, 29), providing a genome-wide (and gene transfer-aware) extension of the classical approach used to root the tree of life on the basis of ancient gene duplications (3, 4).

Our method (amalgamated likelihood estimation, ALE) improves on earlier approaches by explicitly accounting for uncertainty in the gene tree topologies and in the events leading to those topologies, while simultaneously estimating rates of gene duplication, transfer and loss directly from the data (31). Simulations suggest that root inferences under ALE are robust to variation in taxon sampling and the proportion of extinct lineages (fig. S2), that the method finds the correct root even under high levels of gene transfer (1, 29), and that numbers of duplication, transfer and loss events are accurately recovered from the data (figs S3-S8). These results suggest that ALE is appropriate for the problem at hand (36).

Rooting Bacteria without an outgroup

To obtain an unrooted species tree for ALE analysis, we selected a focal dataset of 265 genomes representative of bacterial diversity according to the Genome Taxonomy Database (GTDB)(13). We inferred the tree from a concatenation of 62 conserved single copy markers (table S1) using the LG+C60+R8+F model in IQ-Tree 1.6.10 (Fig. 1), which was chosen as the best-fitting model using the Bayesian Information Criterion (37). This yielded highly congruent trees when removing 20-80% of the most compositionally heterogeneous sites from

the alignment (fig. S9), suggesting that the key features of the topology are not composition-driven LBA artifacts. One exception was the position of the Fusobacteriota, which was recovered as a sister lineage to a clade comprising Deinococcota, Synergistota and Thermotogota (DST) when 20% of the most heterogeneous sites were removed (fig. S2A), but was recovered as a single lineage between Terrabacteria+DST and Gracilicutes in all other trees.

We used ALE to test the support for 62 root positions (tables S3 and S4) on the unrooted topology by reconciling gene trees for 11,272 homologous gene families (inferred using MCL (38)) from the 265 bacterial genomes. Note that this method does not assume that the root lies between Bacteria and Archaea. In addition to testing root positions corresponding to published hypotheses, we exhaustively tested all inner nodes of the tree above the phylum level. The ALE analysis rejected all of the root positions tested (AU test, $P < 0.05$) except for three adjacent branches, lying between the two major clades of Gracilicutes (comprising the majority of diderm lineages) and Terrabacteria (comprising monoderm and atypical diderm lineages; Fig. 1); the difference between the three root positions was the position of the Fusobacteriota in relation to these two major clades (Fig. 1B). Alternative roots were rejected with increasing confidence as distance from the optimal root region increased (Fig. 1C, table S3).

We tested the robustness of the inferred root region by i) excluding gene families with extreme duplication, transfer or loss rates; ii) repeating the analysis using gene families constructed with an assignment to families in the Clusters of Orthologous Genes (COG) (39) ontology, and iii) repeating the analysis on an secondary independent sampling of the tree in which CPR comprise 40% of the genomes (11) (figs S10-S13, table S5). These analyses consistently recovered the root between the Gracilicutes and Terrabacteria, regardless of the position of the Fusobacteriota. A Gracilicutes-Terrabacteria root was previously reported (40, 41) but these studies did not include the CPR, which has recently been suggested to represent the earliest diverging bacterial lineage (11, 16). Our outgroup-free analysis consistently recovered CPR nested within the Terrabacteria, as a sister clade to Chloroflexota+Dormibacterota, even with CPR representing more than 40% of the taxa included. This finding implies that the CPR are the result of reductive evolution from a free-living ancestor, a scenario that has been proposed previously (21).

Transfers contain information about the relative timing of divergences because for each transfer, the donor must be at least as old as the recipient (42, 43). To establish the relative ages of the crown groups of different phyla, we used high-confidence relative age constraints recovered in at least 95/100 bootstrap replicates common to the focal and secondary datasets (36). Simulations suggest that this approach accurately recovers relative clade ages (98.4% accuracy on a simulated dataset the same size as the focal dataset, fig. S14). Our analysis (Fig. 2) predicts that the Firmicutes crown group is the oldest among extant bacterial phyla (median rank 2 ± 1.43 std. dev.) followed by the crown groups of the CPR (median rank 3 ± 2), Proteobacteria (median rank 3 ± 1.59) and Acidobacteriota (median rank 3 ± 1.56), suggesting that these lineages were the earliest to diversify within the Bacteria. The crown groups of lineages predominantly associated with animal hosts, Spirochaetota (median rank 10 ± 0.85) and Elusimicrobiota (median rank 11 ± 0.62), appear to be the youngest among extant phyla.

Is bacterial evolution treelike?

How much of bacterial evolution can be explained by the concept of a rooted species tree? Horizontal gene transfer (HGT) is frequent in prokaryotes, and published analyses indicate that most or all prokaryotic gene families have experienced HGT during their history (1, 44). This implies that there is no single tree that fully describes the evolution of all bacterial genes or genomes (45, 46). Extensive HGT is existentially challenging for concatenation, because it greatly curtails the number of genes that evolve on a single underlying tree (47). Phylogenetic networks (46, 48) were the first methods to explicitly acknowledge non-vertical evolution, but can be difficult to interpret biologically. Gene tree-species tree reconciliation unites tree and network-based approaches by modelling both the horizontal components of genome evolution (a fully reticulated network allowing all possible transfers) and the vertical trace (a common rooted species tree). This framework enables us to quantify the contributions of vertical and horizontal processes to bacterial evolutionary history.

Our analyses (Fig. 3) reveal that most bacterial gene families present in two or more species (9678/10518 MCL families, 92%) have experienced at least one gene transfer during their evolution; only very small families have escaped transfer entirely on the timescales considered here (fig. S15). Consistent with previous analyses (1, 49), transfer rates vary across gene functional categories, with genes encoding proteins involved in defense mechanisms (such as antibiotic biosynthesis) and in the production of secondary metabolites being the most frequently transferred, and those coding for translational and cell cycle proteins the least (Fig. 3B). Despite this accumulation of HGT, most gene families evolve vertically the majority of the time, with mean verticality estimated to be 64% in the focal and 68% in the secondary dataset.

Genome-wide reconciliation of gene trees with the species tree demonstrates that the optimal rooted species tree provides an apt summary of much of bacterial evolutionary history, even for the deepest branches of the tree (50). From the gene's eye view, gene families evolve neither entirely vertically nor horizontally: core genes are occasionally transferred, and even frequently exchanged genes contribute useful vertical signal; for example, the median number of genes that evolve vertically on a branch of the species tree is 998.92 in the focal analysis (table S16), far greater than the number of genes that have been concatenated at the level of all Bacteria. From the perspective of the genome, constituent genes have different ages (or residence times), corresponding to the time at which they originated or were most recently acquired by gene transfer, within the resolution of our taxonomic sampling.

This analysis indicates that, on average, 82% of the genes on all genomes from adequately represented phyla (5 or more genomes) were most recently acquired after the diversification of that phylum, though all genomes retain a smaller proportion (10-27%) of genes that have descended vertically from the stem lineage of their phylum or even earlier (Fig. 3c). There are two explanations for this distribution of gene persistence times: i) *de novo* gene origination within phyla (that is, lineage-specific gene families) and ii) the cumulative impact of gene transfer which curtails gene persistence times when looking back from the present day even though most transmissions are vertical.

Ancestral proteome of the Last Bacterial Common Ancestor (LBCA)

Reconciliation analyses not only allow us to infer the acquisition of genes across the tree, but also to estimate the metabolic potential of the LBCA. We built a second smaller set of COG-based gene families better suited for functional annotation, and reconciled their gene trees with the species tree (36). In the following reconstruction, we indicate when gene content inferences differ between roots (36). PPs for genes directly relevant to our reconstruction are provided in table S7, and all of the pathways we discuss below were confirmed in our analysis of the secondary dataset (36). From the root placement and estimated rates of gene family extinction in the focal analysis (1), we predict that LBCA encoded 1293-2143 COG family members, the majority of which (median estimates 65-69.5%; 95% CI 57-82%) survived to be sampled in at least one present day genome. Based on the relationship between COG family members and genome size for extant Bacteria (Pearson's $r = 0.96$, $P = 8 \times 10^{-153}$), we estimate the genome size of LBCA to be 2.7 ± 0.4 Mb (standard error) for root 1 of the focal analysis (Fusobacteriota with Terrabacteria; Fig. 1(b)); 2.6 ± 0.4 Mb for root 2 (Fusobacteriota with Gracilicutes), and 1.6 ± 0.5 Mb for root 3 (Fusobacteriota root). Under all three roots, the trend in genome size evolution from LBCA to modern taxa is an ongoing moderate increase through time in estimated COG family complements and genome sizes. The most significant departure from this trend is a reduction in genome size of 0.47-0.56 Mb on the CPR stem lineage after divergence from their common ancestor with Chloroflexota+Dormibacterota (fig. S16). COG families lost on the CPR stem include components of the electron transport chain, carbon metabolism, flagellar biosynthesis and motor switch proteins, amino acid biosynthesis, the Clp protease subunit ClpX and RNA polymerase sigma factor-54 (table S8), consistent with their absence in extant CPR (18).

The inferred ancestral gene set for LBCA includes most components of the modern bacterial transcription, translation and DNA replication systems (table S7). It also includes an FtsZ-based cell division machinery and pathways for signal transduction, membrane transport and secretion (Fig. 4, (36)). Further, we identified proteins involved in bacterial phospholipid biosynthesis, suggesting that LBCA had bacterial-type ester-lipid membranes (Fig. 4). We also identified most of the proteins required for flagella and pili synthesis, and for quorum sensing, suggesting that LBCA was motile (51, 52). Since bacterial genes are typically maintained by strong purifying selection (53), these findings imply that LBCA lived in an environment in which dispersal, chemotaxis and surface attachment were advantageous.

Moderate support for the presence of the shape-determining protein *mreB* (0.9/0.88/0.73, root branches 1-3 as depicted in Fig. 1B), *mreC* (0.82/0.79/0.57) and *mreD* (0.86/0.83/0.63) at the root suggests that LBCA were rod-shaped cells (52). We also obtained high root PPs for proteins mediating outer cell envelope biosynthesis including lipopolysaccharides (LPS), from which we infer that LBCA had a double membrane with an LPS layer (36). Consistent with this inference, there was high support for the flagellar subunits FlgH, FlgI and FgA, which anchor flagella in diderm membranes (54), and for the Type IV pilus subunit PilQ, which among extant bacteria is specific to diderms (54, 55). Altogether, this supports hypotheses (9) in which LBCA was a diderm (54–56), and argues against scenarios in which the Gram-negative double membrane originated by endosymbiosis between monoderms (single-membraned bacteria (10)) or via the arrest of sporulation (57) in a spore-forming monoderm ancestor. Thus, diderm-to-monoderm transitions must have occurred subsequently on multiple occasions within Bacteria (54–56).

We recovered components of several core pathways for carbohydrate metabolism with high posterior support, including glycolysis, the tricarboxylic acid (TCA) cycle, and the pentose phosphate pathway (Figs. 4, S17, S18, table S7, (36)). Modern bacteria fix carbon using several different pathways, including the Calvin cycle, the 3-hydroxypropionate bicycle and variations thereof, the reductive glycine pathway (58), the Wood-Ljungdahl Pathway (WLP) and the reverse TCA cycle, the latter two of which have been suggested to have emerged early in the history of life (41, 59–63). Of these, we identified several enzymes of the TCA cycle and the reductive glycine pathway, although we did not recover the key enzymes of either pathway, and the directionality of the recovered enzymes is difficult to assess (64) (Figs. 4, S17 and S18). Furthermore, we identified several enzymes of the methyl branch of the WLP for acetate biosynthesis and components of a putative RNF complex (Figs. 4, S17 and S18), which together may indicate that LBCA was capable of acetogenic growth (65) (36). However, the key enzyme of the WLP, the carbon monoxide dehydrogenase/acetyl-CoA synthase complex (41), had only moderate root support (PP=0.5-0.75) for two subunits and low support (PP <0.5) for other subunits. Thus, while our analyses support the antiquity of components of the WLP, acetogenesis, the TCA cycle and several other core metabolic pathways, they do not confidently establish the combination of pathways employed by LBCA (36).

Finally, our reconstruction also indicated high posterior support for elements of an adaptive immune CRISPR-Cas system (66, 67), including the universally conserved Cas endonuclease, Cas1 (PP=0.96/0.93/0.89), essential for spacer acquisition and insertion into CRISPR cassettes (68, 69). Among other roles, CRISPR systems are crucial in antiviral defense and are activated in response to viral exposure (70); therefore these findings are consistent with hypotheses suggesting that LBCA was already co-evolving with parasitic replicators such as bacteriophages and plasmids (71, 72).

Conclusions

Here, we have used reconciliation methods to model both the vertical and horizontal components of bacterial evolution. These components are complementary, illuminating different facets of bacterial evolution, and we show that the horizontal component can be used to root and orient the vertical tree. Our analyses root the Bacteria between two major clades, the Terrabacteria and Gracilicutes, in contrast to recent outgroup-rooted analyses which place the root on the CPR branch. Instead, we predict that CPR evolved from a common ancestor with the Chloroflexota and Dormibacterota by reductive evolution. We infer that the last bacterial common ancestor was a fully fledged free-living diderm cell with LPS, a multimeric flagellum, and a type III CRISPR-Cas system.

Phylogenetic models are necessarily simplified, and there is much fruitful work to be done to better capture the full heterogeneity of the evolutionary process in the reconciliation framework, from varying diversification rates to endosymbioses. With increased sampling and improved methods, reconciliation analyses should be able to probe still deeper into the early evolutionary history of life on Earth.

Methods (summary)

Phylogenetics: We used two alternative approaches to assemble representative sets of bacterial genomes. In the focal analysis, we sampled 265 genomes evenly from across the GTDB taxonomy (13). In the secondary analysis, we sampled 341 genomes according to the diversity of major bacterial lineages reported in a previous study (11). We used the OMA(73) algorithm to identify candidate single-copy orthologs and manually inspected initial single gene trees to identify a set of 62 congruent phylogenetic markers. Sequences were aligned using mafft 7.453 (74), and trimmed using BMGE 1.12 (75) with the BLOSUM30 matrix. Unrooted species trees were inferred from a concatenation of the 62 markers under the LG+C60+R8+F model in IQ-TREE 1.6.10(76), which was the best-fitting model according to the BIC(37). To perform outgroup rooting analyses, we searched the genomes of 148 Archaea for orthologues of the 62 marker gene set, and identified a subset of 29 genes with congruent single-gene phylogenies. AU tests (28) were performed in IQ-TREE.

Gene tree-species tree reconciliation: To infer gene families, we performed all-versus-all DIAMOND (77) searches among the input protein sets, and clustered the results using the MCL algorithm (38) with an inflation parameter of 1.2. Gene clusters were aligned and trimmed as described above, and bootstrap distributions inferred under the best-fitting model in IQ-TREE. We used ALEml_undated (31) to perform gene tree-species tree reconciliation. The relative ages of bacterial crown groups were estimated with MaxTiC (43) using only those transfer-based age constraints that were recovered in both the focal and secondary datasets. Estimates of gene family and lineage verticality were averaged over the reconciliations obtained in the focal analysis when rooting on each of the three candidate branches in the root region.

Simulations and sensitivity analyses: To evaluate ALE performance, we simulated gene family evolution using Zombi (78) combined with rejection sampling to obtain sets of simulated gene families similar to the real data in terms of inferred DTL events. We then compared simulated and inferred numbers of events under a range of conditions (36). To evaluate the robustness of root inferences, we ordered gene families by decreasing DTL rates, rate ratios, and a range of other proxies for lack of informativeness and potential for introducing bias (36), and compared the likelihoods of competing root hypotheses as increasing proportions of gene families were excluded from the calculation.

Ancestral metabolic reconstruction: We inferred COG gene families by assigning the sequences on each sampled genome to gene families from the COG ontology (39) using eggno-mapper 2 (79), which were then used to perform gene tree-species tree reconciliation. Root origination probabilities for each of the 23 COG functional categories were inferred by maximising the total reconciliation likelihood over all gene families. These category-specific probabilities for origination at the root were then used to estimate the posterior probability (PP) that each gene family was present at the root of the tree. To infer the gene family content and metabolic repertoire of LBCA, functional annotations of protein sequences were obtained and assigned to COG families present at the root. The proteome of LBCA proteome was reconstructed taking into account the respective PPs for key gene families and metabolic pathways.

A detailed account of all analyses is provided in the Supplementary Methods (36).

References

1. T. A. Williams *et al.*, Integrative modeling of gene and genome evolution roots the archaeal tree of life. *Proc. Natl. Acad. Sci. U. S. A.* **114**, E4602–E4611 (2017).
2. J. R. Brown, W. F. Doolittle, Root of the universal tree of life based on ancient aminoacyl-tRNA synthetase gene duplications. *Proc. Natl. Acad. Sci. U. S. A.* **92**, 2441–2445 (1995).
3. J. P. Gogarten *et al.*, Evolution of the vacuolar H⁺-ATPase: implications for the origin of eukaryotes. *Proc. Natl. Acad. Sci. U. S. A.* **86**, 6661–6665 (1989).
4. N. Iwabe, K. Kuma, M. Hasegawa, S. Osawa, T. Miyata, Evolutionary relationship of archaebacteria, eubacteria, and eukaryotes inferred from phylogenetic trees of duplicated genes. *Proc. Natl. Acad. Sci. U. S. A.* **86**, 9355–9359 (1989).
5. O. Zhaxybayeva, P. Lapierre, J. P. Gogarten, Ancient gene duplications and the root(s) of the tree of life. *Protoplasma*. **227**, 53–64 (2005).
6. F. U. Battistuzzi, S. B. Hedges, A major clade of prokaryotes with ancient adaptations to life on land. *Mol. Biol. Evol.* **26**, 335–343 (2009).
7. M. Bocchetta, S. Gribaldo, A. Sanangelantoni, P. Cammarano, Phylogenetic depth of the bacterial genera *Aquifex* and *Thermotoga* inferred from analysis of ribosomal protein, elongation factor, and RNA polymerase subunit sequences. *J. Mol. Evol.* **50**, 366–380 (2000).
8. C. Brochier, H. Philippe, Phylogeny: a non-hyperthermophilic ancestor for bacteria. *Nature*. **417**, 244 (2002).
9. T. Cavalier-Smith, Rooting the tree of life by transition analyses. *Biol. Direct.* **1**, 19 (2006).
10. J. A. Lake, Evidence for an early prokaryotic endosymbiosis. *Nature*. **460**, 967–971 (2009).
11. L. A. Hug *et al.*, A new view of the tree of life. *Nat Microbiol.* **1**, 16048 (2016).
12. S. Mukherjee *et al.*, 1,003 reference genomes of bacterial and archaeal isolates expand coverage of the tree of life. *Nat. Biotechnol.* **35**, 676–683 (2017).
13. D. H. Parks *et al.*, A standardized bacterial taxonomy based on genome phylogeny substantially revises the tree of life. *Nat. Biotechnol.* **36**, 996–1104 (2018).
14. D. H. Parks *et al.*, Recovery of nearly 8,000 metagenome-assembled genomes substantially expands the tree of life. *Nat Microbiol.* **2**, 1533–1542 (2017).
15. C. T. Brown *et al.*, Unusual biology across a group comprising more than 15% of domain Bacteria. *Nature*. **523**, 208–211 (2015).
16. Q. Zhu *et al.*, Phylogenomics of 10,575 genomes reveals evolutionary proximity between domains Bacteria and Archaea. *Nat. Commun.* **10**, 5477 (2019).

17. C. J. Castelle, J. F. Banfield, Major New Microbial Groups Expand Diversity and Alter our Understanding of the Tree of Life. *Cell*. **172**, 1181–1197 (2018).
18. C. J. Castelle, *et al.*, Biosynthetic capacity, metabolic variety and unusual biology in the CPR and DPANN radiations. *Nat. Rev. Microbiol.* **16**, 629–645 (2018).
19. J. P. Beam, *et al.*, "Ancestral absence of electron transport chains in Patescibacteria and DPANN." *Front. Microbiol.* **11**, 1848 (2020).
20. C. J. Castelle *et al.*, Genomic Expansion of Domain Archaea Highlights Roles for Organisms from New Phyla in Anaerobic Carbon Cycling. *Curr. Biol.* **25**, 690–701 (2015).
21. R. Méheust, D. Burstein, C. J. Castelle, J. F. Banfield, The distinction of CPR bacteria from other bacteria based on protein family content. *Nat. Commun.* **10**, 4173 (2019).
22. A. Graybeal, Is it better to add taxa or characters to a difficult phylogenetic problem? *Syst. Biol.* **47**, 9–17 (1998).
23. S. M. Hedtke, T. M. Townsend, D. M. Hillis, Resolution of phylogenetic conflict in large data sets by increased taxon sampling. *Syst. Biol.* **55**, 522–529 (2006).
24. J. Bergsten, A review of long-branch attraction. *Cladistics*. **21** (2005), pp. 163–193.
25. R. Gouy, D. Baurain, H. Philippe, Rooting the tree of life: the phylogenetic jury is still out. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **370**, 20140329 (2015).
26. J. A. Lake, R. G. Skophammer, C. W. Herbold, J. a. Servin, Genome beginnings: rooting the tree of life. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **364**, 2177–2185 (2009).
27. R. G. Skophammer, J. A. Servin, C. W. Herbold, J. A. Lake, Evidence for a gram-positive, eubacterial root of the tree of life. *Mol. Biol. Evol.* **24**, 1761–1768 (2007).
28. H. Shimodaira, An approximately unbiased test of phylogenetic tree selection. *Syst. Biol.* **51**, 492–508 (2002).
29. G. J. Szöllösi, B. Boussau, S. S. Abby, E. Tannier, V. Daubin, Phylogenetic modeling of lateral gene transfer reconstructs the pattern and relative timing of speciations. *Proc. Natl. Acad. Sci. U. S. A.* **109**, 17513–17518 (2012).
30. L. A. David, E. J. Alm, Rapid evolutionary innovation during an Archaeal genetic expansion. *Nature*. **469**, 93–96 (2011).
31. G. J. Szöllösi, W. Rosikiewicz, B. Boussau, E. Tannier, V. Daubin, Efficient exploration of the space of reconciled gene trees. *Syst. Biol.* **62**, 901–912 (2013).
32. L. A. Katz, J. R. Grant, L. W. Parfrey, J. G. Burleigh, Turning the crown upside down: gene tree parsimony roots the eukaryotic tree of life. *Syst. Biol.* **61**, 653–660 (2012).
33. D. M. Emms, S. Kelly, STRIDE: Species Tree Root Inference from Gene Duplication Events. *Mol. Biol. Evol.* **34**, 3267–3278 (2017).
34. A. Zwaenepoel, Y. Van de Peer, Inference of Ancient Whole-Genome Duplications and the Evolution of Gene Duplication and Loss Rates. *Mol. Biol. Evol.* **36**, 1384–1404 (2019).

35. S. Q. Le, O. Gascuel, An improved general amino acid replacement matrix. *Mol. Biol. Evol.* **25**, 1307–1320 (2008).
36. See Supplementary Materials.
37. D. Posada, T. R. Buckley, Model selection and model averaging in phylogenetics: advantages of akaike information criterion and bayesian approaches over likelihood ratio tests. *Syst. Biol.* **53**, 793–808 (2004).
38. A. J. Enright, S. Van Dongen, C. A. Ouzounis, An efficient algorithm for large-scale detection of protein families. *Nucleic Acids Res.* **30**, 1575–1584 (2002).
39. M. Y. Galperin, D. M. Kristensen, K. S. Makarova, Y. I. Wolf, E. V. Koonin, Microbial genome analysis: the COG approach. *Brief. Bioinform.* **20**, 1063–1070 (2019).
40. K. Raymann, C. Brochier-Armanet, S. Gribaldo, The two-domain tree of life is linked to a new root for the Archaea. *P. Natl. Acad. Sci. USA.* **112**, 6670–6675 (2015).
41. P. S. Adam, G. Borrel, S. Gribaldo, Evolutionary history of carbon monoxide dehydrogenase/acetyl-CoA synthase, one of the oldest enzymatic complexes. *Proc. Natl. Acad. Sci. U. S. A.* **115**, E1166–E1173 (2018).
42. A. A. Davín *et al.*, Gene transfers can date the tree of life. *Nature Ecology & Evolution* **2**, 904–909 (2018).
43. C. Chauve *et al.*, <https://www.biorxiv.org/content/10.1101/127548v4> (2017)
44. T. Dagan, W. Martin, Ancestral genome sizes specify the minimum rate of lateral gene transfer during prokaryote evolution. *Proc. Natl. Acad. Sci. U. S. A.* **104**, 870–875 (2007).
45. W. F. Doolittle, Phylogenetic classification and the universal tree. *Science.* **284**, 2124–2129 (1999).
46. W. F. Doolittle, E. Baptiste, Pattern pluralism and the Tree of Life hypothesis. *Proc. Natl. Acad. Sci. U. S. A.* **104**, 2043–2049 (2007).
47. T. Dagan, W. Martin, The tree of one percent. *Genome Biol.* **7**, 118 (2006).
48. D. Alvarez-Ponce, P. Lopez, E. Baptiste, J. O. McInerney, Gene similarity networks provide tools for understanding eukaryote origins and evolution. *Proc. Natl. Acad. Sci. U. S. A.* **110**, E1594–603 (2013).
49. R. Jain, M. C. Rivera, J. A. Lake, Horizontal gene transfer among genomes: the complexity hypothesis. *Proc. Natl. Acad. Sci. U. S. A.* **96**, 3801–3806 (1999).
50. P. Puigbò, Y. I. Wolf, E. V. Koonin, The tree and net components of prokaryote evolution. *Genome Biol. Evol.* **2**, 745–756 (2010).
51. R. Liu, H. Ochman, Stepwise formation of the bacterial flagellar system. *Proc. Natl. Acad. Sci. U. S. A.* **104**, 7116–7121 (2007).
52. F. El Baidouri, C. Venditti, S. Suzuki, A. Meade, S. Humphries, <https://www.biorxiv.org/content/10.1101/2020.08.20.260398v1.full> (2020).
53. I. Sela, Y. I. Wolf, E. V. Koonin, Theory of prokaryotic genome evolution. *Proc. Natl. Acad. Sci. U. S. A.* **113**, 11399–11407 (2016).

54. L. C. Antunes *et al.*, Phylogenomic analysis supports the ancestral presence of LPS-outer membranes in the Firmicutes. *Elife*. **5** (2016).
55. D. Megrian, N. Taib, J. Witwinowski, C. Beloin, S. Gribaldo, One or two membranes? Diderm Firmicutes challenge the Gram-positive/Gram-negative divide. *Mol. Microbiol.* **113**, 659-671 (2020).
56. N. Taib *et al.*, Genome-wide analysis of the Firmicutes illuminates the diderm/monoderm transition. *Nature Ecology & Evolution* **4**, 16661-1672 (2020).
57. E. I. Tocheva, D. R. Ortega, G. J. Jensen, Sporulation, bacterial cell envelopes and the origin of life. *Nat. Rev. Microbiol.* **14**, 535–542 (2016).
58. I. Sánchez-Andrea *et al.*, The reductive glycine pathway allows autotrophic growth of *Desulfovibrio desulfuricans*. *Nat. Commun.* **11**, 5090 (2020).
59. M. C. Weiss *et al.*, The physiology and habitat of the last universal common ancestor. *Nature Microbiology*, **1**, 1–8 (2016).
60. F. L. Sousa, W. F. Martin, Biochemical fossils of the ancient transition from geoeconomics to bioenergetics in prokaryotic one carbon compound metabolism. *Biochim. Biophys. Acta* **1837**, 964-981 (2014).
61. F. L. Sousa, S. Nelson-Sathi, W. F. Martin. One step beyond a ribosome: The ancient anaerobic core. *BBA-Bioenergetics* **1857**, 1027-1038 (2016).
62. G. Borrel, P. S. Adam, S. Gribaldo, Methanogenesis and the Wood-Ljungdahl pathway: an ancient, versatile, and fragile association. *Genome Biol. Evol.* **8**, evw114 (2016).
63. G. Fuchs, Alternative pathways of carbon dioxide fixation: insights into the early evolution of life? *Annu. Rev. Microbiol.* **65**, 631–658 (2011).
64. T. Nunoura *et al.*, A primordial and reversible TCA cycle in a facultatively chemolithoautotrophic thermophile. *Science*. **359**, 559–563 (2018).
65. K. Schuchmann, V. Müller, Autotrophy at the thermodynamic limit of life: a model for energy conservation in acetogenic bacteria. *Nat. Rev. Microbiol.* **12**, 809–821 (2014).
66. K. S. Makarova *et al.*, Evolution and classification of the CRISPR-Cas systems. *Nat. Rev. Microbiol.* **9**, 467–477 (2011).
67. E. V. Koonin, K. S. Makarova, Origins and evolution of CRISPR-Cas systems. *Philosophical Transactions of the Royal Society B: Biological Sciences*. **374** 20180087 (2019).
68. J. K. Nuñez *et al.*, Cas1–Cas2 complex formation mediates spacer acquisition during CRISPR–Cas adaptive immunity. *Nature Structural & Molecular Biology*. **21** 528–534 (2014).
69. K. S. Makarova *et al.*, An updated evolutionary classification of CRISPR-Cas systems. *Nat. Rev. Microbiol.* **13**, 722–736 (2015).
70. R. Barrangou *et al.*, CRISPR provides acquired resistance against viruses in prokaryotes. *Science*. **315**, 1709–1712 (2007).
71. E. V. Koonin, The origins of cellular life. *Antonie Van Leeuwenhoek*. **106**, 27–41

(2014).

72. M. Krupovic, V. V. Dolja, E. V. Koonin, Origin of viruses: primordial replicators recruiting capsids from hosts. *Nat. Rev. Microbiol.* **17**, 449–458 (2019).

73. A. C. J. Roth, G. H. Gonnet, C. Dessimoz, Algorithm of OMA for large-scale orthology inference. *BMC Bioinformatics.* **9**, 518 (2008).

74. K. Katoh, D. M. Standley, MAFFT multiple sequence alignment software version 7: improvements in performance and usability. *Mol. Biol. Evol.* **30**, 772–780 (2013).

75. A. Criscuolo, S. Gribaldo, BMGE (Block Mapping and Gathering with Entropy): a new software for selection of phylogenetic informative regions from multiple sequence alignments. *BMC Evol. Biol.* **10**, 210 (2010).

76. L. T. Nguyen, H. A. Schmidt, A. Von Haeseler, B. Q. Minh, IQ-TREE: A fast and effective stochastic algorithm for estimating maximum-likelihood phylogenies. *Mol. Biol. Evol.* **32**, 268–274 (2015).

77. B. Buchfink, C. Xie, D. H. Huson, Fast and sensitive protein alignment using DIAMOND. *Nat. Methods.* **12**, 59–60 (2015).

78. A. A. Davín, T. Tricou, E. Tannier, D. M. de Vienne, G. J. Szöllősi, Zombi: a phylogenetic simulator of trees, genomes and sequences that accounts for dead lineages. *Bioinformatics.* **36**, 1286–1288 (2020).

79. J. Huerta-Cepas *et al.*, Fast Genome-Wide Functional Annotation through Orthology Assignment by eggNOG-Mapper. *Mol. Biol. Evol.* **34**, 2115–2122 (2017).

80. Online Data Supplement doi:10.6084/m9.figshare.12651074.

81. N. Dombrowski *et al.*, Undinarchaeota illuminate DPANN phylogeny and the impact of gene transfer on archaeal evolution. *Nat Commun* **11**, 3939 (2020).

82. M. Wilkinson, J. O. McInerney, R. P. Hirt, P. G. Foster, T. M. Embley, Of clades and clans: terms for phylogenetic relationships in unrooted trees. *Trends Ecol. Evol.* **22**, 114–115 (2007).

83. S. M. van Dongen, *Graph Clustering by Flow Simulation*, thesis, University of Utrecht, Utrecht, NL (2000).

84. D. H. Parks, M. Imelfort, C. T. Skennerton, P. Hugenholtz, G. W. Tyson, CheckM: assessing the quality of microbial genomes recovered from isolates, single cells, and metagenomes. *Genome Res.* **25**, 1043–1055 (2015).

85. G. J. Szöllősi, E. Tannier, N. Lartillot, V. Daubin, Lateral gene transfer from the dead. *Syst. Biol.* **62**, 386–397 (2013).

86. H.-C. Wang, B. Q. Minh, E. Susko, A. J. Roger, Modeling Site Heterogeneity with Posterior Mean Site Frequency Profiles Accelerates Accurate Phylogenomic Estimation. *Syst. Biol.* **67**, 216–235 (2018).

87. E. Susko, A. J. Roger, On reduced amino acid alphabets for phylogenetic inference. *Mol. Biol. Evol.* **24**, 2139–2150 (2007).

88. C. Zhang, E. Sayyari, S. Mirarab, ASTRAL-III: increased scalability and impacts of contracting low support branches in *Comparative Genomics* (Springer, Cham, 2017),

Lecture Notes in Computer Science, pp. 53–75.

89. S. R. Eddy, Accelerated profile HMM searches. *PLoS Comput Biol* **7**, e1002195 (2011).
90. R. D. Finn, J. Clements, S. R. Eddy, HMMER web server: interactive sequence similarity searching. *Nucleic Acids Res.* **39**, W29–37 (2011).
91. T. Aramaki *et al.*, KofamKOALA: KEGG Ortholog assignment based on profile HMM and adaptive score threshold. *Bioinformatics.* **36**, 2251–2252 (2020).
92. P. Jones *et al.*, InterProScan 5: genome-scale protein function classification. *Bioinformatics.* **30**, 1236–1240 (2014).
93. J. Huerta-Cepas *et al.*, eggNOG 5.0: a hierarchical, functionally and phylogenetically annotated orthology resource based on 5090 organisms and 2502 viruses. *Nucleic Acids Res.* **47**, D309–D314 (2019).
94. J. Huerta-Cepas *et al.*, eggNOG 4.5: a hierarchical orthology framework with improved functional annotations for eukaryotic, prokaryotic and viral sequences. *Nucleic Acids Res.* **44**, D286–93 (2016).
95. R. L. Tatusov *et al.*, The COG database: an updated version includes eukaryotes. *BMC Bioinformatics.* **4**, 41 (2003).

Acknowledgements

Funding: GAC is supported by a Royal Society Research Grant to TAW. TAW is supported by a Royal Society University Research Fellowship and NERC grant NE/P00251X/1. GJSz received funding from the European Research Council under the European Union’s Horizon 2020 research and innovation program under Grant Agreement 714774 and Grant GINOP-2.3.2.-15-2016- 00057. AS is supported by the Swedish Research Council (VR starting grant 2016-03559 to AS) and the NWO-I foundation of the Netherlands Organisation for Scientific Research (WISE fellowship to AS). AAD and PH were supported by an Australian Research Council Laureate Fellowship (grant no. FL150100038). **Author contributions:** The project was conceived by TAW, GJSz, PH, AS, GAC and AAD. GAC, AAD, TAW, LSz and GSz performed phylogenomic analyses. GJSz developed new analytical methods. TAM, GAC, AAD and AS performed metabolic annotations and reconstructions. All authors contributed to interpretation and writing. **Competing interests:** The authors have no competing interests to declare. **Data and code availability:** All data and code are provided online at (80). New methods are described in detail in the Supplementary Methods (36).

Contents of the Supplementary Materials

Supplementary Materials and Methods

Figs. S1-19

Tables S1-13

References (81-95)

Figure Legends

Fig. 1: A rooted phylogeny of Bacteria. (a) We used gene tree-species tree reconciliation to infer the root of the bacterial tree. The unrooted maximum likelihood phylogeny was inferred from a concatenation of 62 marker genes under the best-fitting model, LG+C60+R8+F, which accounts for site-heterogeneity in the substitution process and uses a mixture of eight substitution rates estimated from the data to model across-site evolutionary rate variation. Branches are coloured according to bootstrap support value. The root falls between two major clades of Bacteria, the Gracilicutes and the Terrabacteria, on one of three statistically equivalent adjacent branches indicated by arrows, shown as rooted trees in (b). All alternative roots tested were rejected (tables S3 and S4) with likelihoods decreasing as a function of distance from the root region, as shown in (c). Previously proposed root positions, including the CPR root, are highlighted in red. FCB are the Fibrobacterota, Chlorobia, Bacteroidota and related lineages; PVC are the Planctomycetota, Verrucomicrobiota, Chlamydiia and related lineages; DST are the Deinococcota, Synergistota and Thermotogota; ACD are Aquificota, Campylobacterota, and Deferribacterota; FA are Firmicutes and Actinobacteriota.

Fig. 2: Relative crown group ages of major bacterial phyla. Gene transfers that occurred during the diversification of Bacteria provide a record of the temporal sequence of events. We used the information provided by directional (donor-to-recipient) patterns of gene transfer to infer the relative ages of bacterial crown groups, focusing on phyla represented by at least 5 genomes in both of our datasets. To summarise this time information, we sampled 1000 time orders that were fully compatible with the constraints recovered from both datasets. (a) Pairwise relative ages of phyla: the proportion of sampled time orders in which each phylum on the x-axis was recovered as younger than each phylum on the y-axis. (b) Relative age distributions of major phyla: for each sampled time order, we ranked the phyla from oldest (1) to youngest (11), and plotted the distribution of the ranks. The crown group radiations of Firmicutes, CPR, Proteobacteria and Acidobacteriota appear to be the oldest among sampled phyla, while those of Elusimicrobiota and Spirochaetota are the youngest.

Fig. 3: The verticality of bacterial genome evolution. (a) The rooted bacterial species tree (Fig. 1) with branches colored according to verticality: the fraction of genes at the bottom of a branch that descend vertically from the top of that branch (see inset; V = vertical, O = origination, T = transfer into a branch; (36)). Node heights reflect relative time order consistent with highly-supported gene transfers (Fig. 2). (b) Transfer propensity by COG functional category: that is, the proportion of gene tree branches that are horizontal $T/(V+T)$ for COG gene families. Genes involved in information processing, particularly translation (J), show the lowest transfer propensity (median 0.31), while genes involved in cell defense mechanisms (V, such as genes involved in antibiotic defense and biosynthesis) are most frequently transferred (median transfer propensity 0.47). (c) From the genome's eye view, this combination of vertical and horizontal processes gives rise to a distribution of gene persistences (residence times), reflecting the point in evolutionary history (within the Crown group, on the Stem, or earlier --- Before) at which the gene was most recently acquired. Across all phyla examined, 82% of genes on sampled genomes were most recently acquired since the crown group radiation of that phylum. Lineage acronyms are as per Figure 1.

Fig. 4: Ancestral reconstruction of the last bacterial common ancestor (LBCA). The reconstruction is based on genes that could be mapped to at least one branch within the root region with a PP >0.5 (figs. S17 and S18, (36)). The presence of a gene within a pathway is indicated as shown in the key. Our analyses suggest that LBCA was a rod-shaped, motile, flagellated double-membraned cell. We recover strong support for central carbon pathways, including glycolysis, the TCA cycle and the pentose phosphate pathway. We did not find unequivocal evidence for the presence of a carbon fixation pathway, although we found moderate support for components of both the Wood-Ljungdahl pathway and the reverse TCA cycle. Though not depicted here, our analyses suggest that the machinery for core information processing and quorum sensing was also present in LBCA (table S7).

a

b

c

a

b

c

Supplementary Materials for

A Rooted Phylogeny Resolves Early Bacterial Evolution

Gareth A. Coleman*, Adrián A. Davín*, Tara A. Mahendrarajah, Lénárd L. Szánthó, Anja Spang, Philip Hugenholtz*, Gergely J. Szöllősi*, Tom A. Williams*

*Contributed equally to this paper.

*Co-corresponding authors; contributed equally to this paper; correspondence to p.hugenholtz@uq.edu.au, ssolo@elte.hu, tom.a.williams@bristol.ac.uk

This PDF file includes:

Materials and Methods

Figs. S1 to S18

Tables S1 to S13

Materials and Methods

Taxon sampling for focal analysis

To obtain a representative taxon sampling from across known bacterial diversity, we sampled taxa according to the classification provided by the Genome Taxonomy Database (GTDB r89)(13) as follows. First, we removed genomes with Quality < 0.75 (Quality is defined as Completeness - (5*Contamination)(14)), and filtered out all phyla subsequently left with fewer than 10 species. Genomes were sampled from the remaining taxa on a per-class basis: for classes containing a single order, the genome with the highest quality score was sampled; for classes containing multiple orders, the highest quality genome from each of two randomly chosen orders was sampled. This protocol ensured that every class in the GTDB is represented in the final tree. We then manually added the genome of *Gloeomargarita litophora* given its importance in constraining the phylogeny and timing of chloroplast evolution. The list of genomes can be found in table S9.

Unrooted species tree inference

We used Orthologous Matrix (OMA) 2.1.1 (73) to identify candidate single-copy bacterial orthologs, and retained those with at least 75% of all species represented in each family. Sequences were aligned in MAFFT (74) using the -auto option, and trimmed in BMGE 1.12 (75) using the BLOSUM30 model. Initial trees were inferred for each candidate marker gene under the LG+G+F model in IQ-TREE 1.6.10 (76). The trees were manually inspected, and we selected orthologues where the monophyly of 14 pre-defined major lineages was not violated with bootstrap support >70%, resulting in 62 final orthologues. Concatenation of this marker set resulted in an alignment of 18,234 amino acids. We inferred an unrooted phylogeny from this concatenate under the LG+C60+R8+F model, which was chosen as the best-fitting model by the BIC criterion in IQ-TREE (76). We additionally removed the most compositionally heterogeneous sites from the sequence alignment using Alignment Pruner (81) (<https://github.com/novigit/davinciCode/blob/master/perl>) (20%, 40%, 60% and 80% respectively) and inferred trees using the same procedure described above in order to compare the resulting topologies.

Outgroup rooting

To root the bacterial tree using an archaeal outgroup, we used a representative sampling of 148 archaeal genomes and inferred the ML tree in IQ-TREE under the best-fitting LG+C60+R8+F model. The concatenated alignment included a subset of 29 out of the 62 bacterial orthologs that were shared between bacteria and archaea, as determined by by Hidden Markov Model (HMM) searches and manual inspection of single gene trees. ML trees for these 29 genes recovered the cladehood (82) of Bacteria and Archaea. We performed approximately-unbiased (AU) tests (28) to determine whether a range of published alternative rooting hypotheses (table S2) could be rejected, given the model and data (AU p-value > 0.05).

Gene family clustering and ALE analysis

We used the protein annotations provided by GTDB, which were originally obtained using Prodigal. To infer homologous gene families for amalgamated likelihood estimation (ALE), we performed an all vs all similarity search using Diamond (77) with an E-value threshold of $<10^{-7}$ to avoid distant hits and $k = 0$ to report all the relevant hits. Current clustering methods are not consummate and the parameters that determine the granularity of clustering do not have a direct biological motivation. Setting the value of the Markov Cluster (MCL) algorithm (83) inflation parameter therefore involves a trade-off between inferring large, inclusive clusters that will contain false positives (sequences that are not part of the real gene family) and small, conservative clusters that may divide real gene families into several subclusters. An additional practical concern for phylogenomics is that overly large clusters may align poorly and result in low-quality single protein trees. In our rooting analysis, we experimented with a range of values for the mcl inflation parameter, and chose 1.2 because the clusters were inclusive without a substantial reduction in post-masking alignment length compared to more granular settings.

Clustering using MCL (83) with an inflation parameter of 1.2 resulted in 186,827 gene families and a total of 11,765 families with 4 or more sequences. We aligned the 11,765 gene families using MAFFT (74) (with the --auto option) and filtered with BMGE (75) (using `bmge -t AA -m BLOSUM30`). After filtering, 260 alignments contained no high-quality columns and were discarded. We filtered out sequences comprising more than 80% gaps to produce the final set of alignments. We also discarded all alignments with less than 30 columns, leaving a total of 11,272 families. The gene trees were computed using IQ-TREE v 1.6.10 using the following command:

```
iqtree -m TEST -s FAMXXX.faa.aln.trimmed -bb 10000 -wbtl -nt AUTO -madd LG4X, LG4M, LG+C10, LG+C20, LG+C30, LG+C40, LG+C50, LG+C60, C10, C20, C30, C40, C50, C60.
```

Conditional clade probabilities (CCPs) were computed using ALEobserve and the resulting ALE files were reconciled with the species tree. Loss rates were corrected by genome completeness, estimated using CheckM (84). We tested 62 roots (Online Data Supplement (80)).

Simulations to evaluate the performance of ALE

Performance of ALE for species tree rooting

The ability of the ALEml_undated algorithm to infer the correct gene tree root in the presence of gene duplications, transfers and losses was previously investigated using simulations (1). Briefly, gene families were simulated on a rooted species tree using a continuous-time origination, duplication, transfer and loss (ODTL) process (that is, a more complex model of genome evolution than that implemented in ALEml_undated), and ALEml_undated was used to estimate the root from subsamples of the simulated families. The maximum likelihood root according to ALE was the correct root in 95/100 replicates, and the log likelihood of alternative roots decreased with nodal distance from the correct root (as observed in our empirical data, see Fig. 1). In the remaining 5 cases, the maximum likelihood root

was one branch away from the true root. Analysis of empirical data suggested that ALE root inferences are robust to (that is, consistent across) subsets of the data that vary in terms of the rate of horizontal gene transfer or species representation in gene families (1). These properties make the ALE approach appropriate for inferring the root of Bacteria.

Accuracy of duplication, transfer and loss rates inferred by ALE

To test the accuracy of ALE at correctly inferring duplication, transfer and losses, we simulated a species tree of 265 leaves (the same size as the focal dataset) using Zombi (78). A simulation approach is necessary because, for empirical data, we do not know the true gene family history and so cannot evaluate method performance directly. The empirical realism of simulations is often an issue, and it is not always clear how best to accommodate the complexities of real genome evolution, including heterogeneity of DTL rates across families and, indeed, biases in rates (for instance, the high lineage-specific rate of gene loss that appears to characterise CPR). To make our simulations as realistic as possible, we first simulated the evolution of gene families using the Gm mode in Zombi, which assumes that every family has its own and independent rates of D, T and L, with the rates sampled from distinct gamma distributions ($D \sim G(0.2, 0.5)$, $T \sim G(2, 0.5)$, $L \sim G(2.2, 0.5)$). We simulated a total of 97929 families. We computed gene tree-species tree reconciliations with ALEml_undated for all families. Then, to ensure that the simulated dataset was as similar as possible to the real dataset, we sampled 2000 families at random from the real dataset. For each of those families, we selected the simulated family most similar in terms of DTL events (similarity was computed as one minus the squared sum of differences between the different inferred events by the reconciliations and the size of the families). This procedure resulted in a set of simulated families that were closely similar to the real data in terms of DTL events (fig. S3). These families recapitulate the gene family- and lineage-specific heterogeneity of DTL rates observed in the real data, and so provide the best possible basis for evaluating the performance of ALE. We used this set of simulated families in the analyses described below.

Comparison of the real (that is, simulated) and inferred numbers of D, T and L events on these data suggest that ALE accurately estimates the numbers of all three kinds of events (fig. S4), with mean errors close to 0 for all three types of events ($D \sim 0.005$, $T \sim 0.048$, $L \sim 0.019$). A detailed examination of the errors that do occur indicated that errors are most common in small families (fig. S5), and that the number of DTL events tends to be under-estimated when the true number of events is high (fig. S6). These observations motivate some of the sensitivity analyses of the empirical data described below, in which gene families with high inferred rates and, in a separate analysis, small sizes were excluded from the root calculation (see “Testing the robustness of the inferred root region” below); the Gracilicutes-Terrabacteria root was robust to all of these treatments.

Different combinations of DTL events can give rise to the same gene tree topology. For example, genes that are patchily distributed across species might be explained by a series of gene transfers, ancestral presence followed by independent losses, or a combination of processes. To investigate whether ALE can distinguish between

different kinds of DTL events based on gene tree topologies, we examined the correlations in inference errors for different kinds of events. Negative correlations (for example, over-estimation of transfer associated with under-estimation of losses) would suggest that the method can mistake one kind of DTL event for another. No correlations of this type were obtained (fig. S7), suggesting that ALE can distinguish the history of DTL events giving rise to a given gene tree topology. To further investigate whether ALE overestimates the number of gene transfers compared to duplications and losses, we specifically examined the inference results for the subset of simulated families with 0 transfers but one or more duplication and loss events (2429 of 97292 families). Of these 2,429 families, ALE correctly inferred that 2,332 (96%) had no transfers.

Next, we investigated whether biases in DTL rates across the tree - which result in variation in genome sizes, such as the small genomes of CPR - impact ALE inference accuracy. To do so, we performed two additional simulations: one in which gene family originations occur at random on the tree (which results in homogeneous simulated gene contents and genome sizes), and one in which gene originations were constrained to occur at the same points as they do in the real data. This latter simulation results in data that recapitulate the variation in gene content and genome size observed in the empirical data. We then compared inference accuracy on the two datasets (Fig. S8). The results suggest that genome size heterogeneity does not substantially affect inference accuracy, with all errors centred on 0 in both datasets (mean errors without heterogeneity: D $\sim$ 0.00342, T $\sim$ 0.0044, L $\sim$ 0.0238; with heterogeneity: D $\sim$ 0.00349, T $\sim$ -0.0007, L $\sim$ -0.0268).

Finally, we investigated the impact of lineage extinction on the accuracy of DTL estimates. In principle, ALE estimates ought to be robust to lineage extinction because gene acquisitions from extinct (or unsampled) lineages are accounted for in the method (85). To investigate, we used Zombi (78) to simulate 1000 species trees with 30 extant (sampled) taxa, with the speciation rate equal to the extinction rate. We performed simulations on species trees with relatively small numbers of tips because, since Zombi is a forward simulator, simulating trees with 265-341 tips in the context of high extinction rates is computationally intractable (that is, only a very small proportion of simulated trees will grow to have 100s of extant tips when the extinction rate is equal to or larger than the speciation rate). 100 gene families were simulated on each species tree; of these, 69921 had at least 4 surviving gene copies in the extant tip genomes and could be used for comparison of DTL inference accuracy. For each family, we calculated inference accuracy as (inferred number of events - simulated number of events)/family size, as above, and evaluated the relationship between the proportion of extinct lineages on the species tree and inference accuracy (fig. S2). We detected a statistically significant but quantitatively small impact of extinction on accuracy, with a higher proportion of extinct lineages corresponding to a slight increase in error (Correlation coefficients between accuracy and proportion of extinct lineages: -0.024 (D); 0.041 (T); -0.093 (L); fig. S2); errors remain centred on 0 even when almost all lineages had gone extinct.

Testing the robustness of the inferred root region

Simulations (see above) are the most direct way to evaluate the performance of ALE, because they provide a controlled situation in which we know the truth with certainty. However, real data are heterogeneous in ways that are difficult to recapitulate in simulations, particularly in terms of variation and biases in the rates of evolutionary processes (DTL, speciation, extinction and substitution rates) across the tree. We therefore performed a range of sensitivity analyses to evaluate the robustness of the inferred root region.

Distributions of DTL rates, rate ratios, and the impact of excluding gene families from the root calculation based on these and other criteria.

Firstly, we ranked gene families by inferred duplication, transfer and loss rates and rate ratios (figs S10-11), and performed a gene-filtering analysis (fig. S12) in which families at either end of the distribution were progressively removed and root likelihoods re-evaluated. This approach is analogous to fast site removal, in that families with very high or low rates may be difficult to model and so mislead inference.

Based on these rankings, we performed gene filtering analyses in which the highest ranked families were progressively removed and the difference in likelihoods between roots (ΔLL) re-evaluated (fig. S12A-S). The first of these plots (fig. S12B-C) illustrates the effect of progressively filtering out the most widely-distributed families (the metric is the number of species with at least one gene in the family). For each pair of plots, the bottom panel shows the threshold value corresponding to the percentile removed. For example, removing 5% of the families represented in the most species corresponds to a threshold of being represented in 150 species or more for MCL families. Note that these threshold plots can also be interpreted as the cumulative distribution of the ranking criteria, i.e. the above example implies that 5% of MCL families are represented on 150 or more genomes. The left hand side of each plot indicates the summed likelihood of each candidate root position (Fig. S12A) on all of the data; moving to the right along the x-axis illustrates how the summed likelihood for each root changes as families at the top end of the distribution are filtered out. Filtering out broadly-distributed families has a similar effect in both the MCL and Clusters of Orthologous Genes (COG) datasets: ΔLL starts to diminish. That is, ALE starts to lose the ability to distinguish between different root hypotheses. Conversely, when ranking families by the opposite criterion (the number of species absent, fig. S12D-E), the difference in likelihood between the roots remains unchanged until a very large fraction of families is excluded.

We next evaluated whether families with different estimated D, T or L rates agree on the optimal root region (that is, whether the root signal from families with high, moderate and low rates is consistent). To evaluate the signal in a root-independent manner, we calculated the mean per family D, T and L rates weighted according to the likelihood of that family for different roots. We can see in fig 12F-K that removing even a substantial (10-20%) fraction of families with the highest D, T or L rates does not change the order of likelihoods for different roots, and it is only after nearly half the families are discarded that we start to see a loss of resolution.

To filter families with potentially problematic rates, we also calculated each of the 6 possible rate ratios (D/T, T/D, D/L, L/D, T/L and L/T) and ranked families according to the maximum of these six ratios (fig. S12L-M). While our results remain unchanged as long as fewer than 10% of families with the most extreme rate ratios are removed, removing between 10 and 30% of the families with the most extreme rate ratios changed the order of the most likely roots while still retaining resolution (see for colors). This indicates that the Fusobacteriota root (dark green) derives its support, in part, from families that are outliers in terms of DTL rate ratios.

Performing further analogous threshold analyses (fig. S12N-S) in terms of mean copy number and verticality, we find that the rooting analysis is robust to filtering based on these criteria.

Effect of gene family clustering

We next evaluated the impact of gene family clustering on our analysis by repeating the entire analysis using COG (39) families; the same root region was recovered with the addition of one adjacent branch, with the reduced resolution likely due to the smaller size of the COG family dataset (3,723 vs 11,272 mcl families; see table S5).

Effect of taxon sampling

Taxon sampling is known to be an important factor in phylogenetic analyses (22), so we next evaluated whether our method of sampling taxa representatively from GTDB impacted our results. To do so, we repeated the entire analysis using a different approach, selecting representative taxa from major bacterial clades previously described in the literature (11,15). We based our sampling on the analysis of (11), in which CPR comprise 40% of bacterial diversity. To do so, we inferred a tree of the bacterial portion of the published (11) concatenate under the LG+G4+F model in IQ-TREE. We divided the tree into 7 major bacterial clades based on a literature search (table S12) and additional environmental lineages with branch length diversity comparable to the known groups. For each group defined in this way, we manually subsampled taxa so as to maintain genetic diversity, while avoiding the longest and shortest branches. We sampled 341 species, comprising 200 ‘classic’ bacteria, 124 CPR bacteria and one bacterial genome respectively from each of the 17 new phyla described by (14); see table S9. We used the same marker gene set as in the focal analysis. A species tree was inferred in IQ-TREE using the LG+C20+G4 model with PMSF (86). Additional trees were inferred in PhyloBayes under the CAT+GTR+G4 model using a recoded alignment using the four-category scheme of Susko and Roger (87), and under the multispecies coalescent model in ASTRAL (88). To infer homologous gene families, we used the same pipeline as that used in the focal analysis. This resulted in 11,781 gene families with 4 or more sequences, which were analysed as in the focal analysis. The unrooted species trees are congruent with each other, except in the placement of a small group of phyla comprising Fusobacteriota, Aquificota, Synergistota, Spirochaetota and Thermotogota (“FASST”). These taxa are resolved in different positions in each of the three unrooted topologies, and are found to be monophyletic in the tree inferred from the recoded alignment. Similarly to the focal analysis, the ALE analyses yielded two

root positions that could not be rejected (AU test, $p > 0.05$), summarised in fig. S13. Both of these rooted phylogenies are congruent with that of the focal (GTDB) analysis, with Terrabacteria and Gracilicutes on either side of the root, with the only differences being in the placement of the FASST taxa. In the focal analysis, as well as the LG+PMSF+G4 and ASTRAL trees inferred as part of the GTDB-independent analysis, FASST were not recovered as a monophyletic group.

Inference of relative divergence times of bacterial clades

We applied the pipeline documented and implemented in the Online Data Supplement (80, RelativeDating.zip) to obtain the relative dated trees. We performed the analysis 5 times: 3 times in the focal dataset, and 2 times in the secondary dataset (using all the roots that could not be rejected). The pipeline consists of parsing the transfers inferred by ALEml_undated (using the MCL families) and discarding those with posterior probability < 0.05 . We used bootstrapping to estimate constraint support in the following way: for each of the three branches in the root region, we sampled the gene families 100 times with replacement and, for each replicate, converted detected transfers to constraints and performed a Maximum Time Consistency (MaxTiC) analysis (42, 43). We then selected the phyla that were represented by 5 genomes or more in both datasets. We pruned the tree of the focal dataset to maintain only those phyla. Finally, we selected the constraints that were supported by both datasets and use those constraints to generate 1000 time orders compatible with those constraints the script `order_explorer.py` (80)). We then ranked all interior nodes on the tree, with the root node having rank 0 and the most recent speciation node having rank 11.

Quantifying vertical and horizontal signals in bacterial genome evolution

In the context of our analyses, “verticality” is the proportion of inferred evolutionary events on a branch of the rooted species tree that reflect vertical descent, estimated using gene tree-species tree reconciliation. We defined branch-wise verticality as $V/(V+O+T)$, where V is the inferred number of vertical transmissions of a gene from the ancestral to descendant ends of the branch; O is the number of new gene originations on the branch; and T is the number of gene transfers into the branch. We defined transfer propensity as $T/(V+T)$, where V and T refer to inferred numbers of events within the history of a gene family (table S11). The numbers reported in the main text have been averaged over the reconciliations obtained using the three possible roots of the focal analysis.

Ancestral gene content and metabolic reconstruction

Protein and protein family functional annotation

Protein sequences from all genomes used for phylogenetic analyses in this study were annotated using a variety of databases. Functional annotations were obtained using `hmmsearch v3.1b2` (settings: `-E 1e-5`) (89,90) against KEGG Orthology (KO) annotations from the KEGG Automatic Annotation Server (KAAS; downloaded April 2019) (91). Additionally, all proteins were scanned for protein domains using `InterProScan (v5.31-70.0; settings: --iprlookup --goterms)` (92).

Multiple hits corresponding to the individual domains of a protein are reported using a custom script (`parse_IPRdomains_vs2_GO_2.py`). For the functional annotation of the 4256 COG families investigated in our ancestral reconstructions, we assigned KOs using a majority rule: i.e. we assigned the KO, that was reported in > 50% of the sequences comprising each of the COG families yielding a COG-to-KO mapping file. Subsequently, we mapped COG descriptions, COG Process/Class, Category description, kegg id, kegg description, and kegg pathway to the COG-to-KO mapping file. COG descriptions were collected from the root annotations (`1_annotations.tsv`) downloaded at EggNOG (v5.0.0) (93). COG functional category and Process/Class descriptions were derived from eggNOG (v4.0) (94). KO pathways were manually curated based on an in-house KO-to-pathway mapping file, and were subsequently mapped to the respective KO. The scripts for annotation and mapping are included in the Online Data Supplement (80).

COG gene families for ancestral gene content reconstruction

We built a set of gene families based on the COG (95) database for ancestral functional inference. To do so, we annotated each genome in the dataset using eggNOG-mapper v2(79), then clustered proteins into families based on their COG annotations. For proteins annotated with more than one COG category (8% of proteins), we included the protein in both COG families. This resulted in 4256 COG families, of which 3723 had 4 or more sequences. COG families are ideal for ancestral reconstruction because they comprise all of the sequences on extant genomes that can be annotated with a given unambiguous function from the COG ontology. In addition, the hierarchical nature of the COG classification (comprising gene family annotations nested within 23 broader functional categories) enabled us to explicitly model the different evolutionary ages of gene functional classes as part of the analysis, by using category-specific root origination priors (see below).

Our COG families are useful for functional reconstruction, but are perhaps less well suited for investigating other aspects of bacterial evolution because they are constructed only from proteins that could be annotated with eggNOG-mapper. By contrast, MCL families represent --- within the limitations of the clustering approach, as discussed above --- an unbiased view of gene family diversity for the set of genomes we analyzed. We therefore base analyses other than those regarding the functional annotation of LBCA on the MCL families. However, since gene clustering methods are not consummate and each has strengths and weaknesses, we also investigated the root signal from the COG families. This analysis identified a similar root region of four adjacent branches, comprising the root region from the focal analysis (3 branches) plus one additional root, in which Spirochaetota branched on the Terrabacteria side of the root (table S5).

Root gene mapping approach

To estimate root presence posterior probabilities (PPs) for each gene family for each of the three supported roots, we first estimated the probability of origination at the root (O_R) by maximum likelihood, finding the O_R value that maximises the total reconciliation likelihood summed over all gene families (table S11). We then used the global ML O_R value to calculate the root presence posterior probabilities

for each family; that is, the probability that one or more copies of a given gene family were present at the root, given the ML O_R value. We estimated root origination rates independently for each of the 23 COG functional categories, and used these rates to estimate the posterior probability of presence at the root node for each gene family. Note that, for all nodes of the tree (including the root nodes), we additionally estimated PPs directly from the sampled reconciliations. Python code implementing this procedure is provided in (80) at Code/O_R_Optimization.py.

Initial gene content and metabolic inferences at a particular node were based on gene families with a posterior presence probability (PP) of >0.95 at that node. This approach is conservative and could miss the presence of certain pathways which may be represented by proteins with a range of PP values. Therefore, we manually investigated the PPs of pathways discussed in this manuscript and inferred the presence of specific pathways or functional modules if the majority of the components were found with PP >0.50, as described in the main text, Fig. 4 and fig. S17.

Impact of root branch on LBCA gene content

The credible set of root branches from the ALE analysis comprised three adjacent branches at the centre of the tree (Fig. 1b). The difference between these three root positions relates to the placement of Fusobacteria, either as the root branch or as the most basal split on either the Gracilicutes or Terrabacteria+DST “sides” of the rooted tree. We therefore estimated root PPs for COG families on all three branches; root PPs under all three roots are provided in the Online Data Supplement (80).

Metabolic comparisons

Results from the PP analysis were used as the framework for metabolic comparisons and reconstruction of the proteome of LBCA. First, the occurrence of an individual COG family across each taxon was counted in R (v3.6.3) (table S4). This binary presence/absence matrix was combined with the PP values for Nodes corresponding to the CPR, Chloroflexota+CPR, Chloroflexota, Terrabacteria, DST+Terrabacteria, Gracilicutes-Spirochaetota, Gracilicutes+Spirochaetota, Root 1, Root 2, and Root 3, filtered with a cutoff of PP>0.50. The combined count table was summarized using the ddply function of the plyr package (v1.8.4), which was used to summarize the counts across each phylogenetic cluster, node, and root. Data is visualized in a heatmap generated using the ggplot function with geom_tile and facet_grid of the ggplot2 package (v3.2.0). Heatmap categories for pathways were scaled based on the number of COG families, results were plotted using the grid.draw function of the grid package (v3.6.3). Heatmaps were manually merged with a species tree in Adobe Illustrator (v22.0.1).

Evaluating the robustness of the LBCA reconstruction to taxon sampling based on the secondary dataset

To evaluate the impact of taxon sampling, we analyzed both the primary and secondary datasets with the same ancestral reconstruction pipeline. Across all

protein families, agreement between the datasets was highly significant, though lower than among the root regions for each dataset (table S13). The PPs between the focal and secondary analysis were significantly correlated independent of rooting (Pearson's correlation 0.48-0.53 $p < 10^{-16}$, see table S13). PPs between root positions in the context of the same analysis were very strongly correlated, with the highest correlation between root 1 and root 2 of the secondary analysis (0.96) and roots 1 and 2 of the primary analysis (0.94). Overall, root 3 of the focal analysis (corresponding to the Fusobacteriota root) correlated the least with other roots. The focal analysis compared to the secondary analysis considered fewer taxa (265 vs 341) and correspondingly fewer COGs (3782 vs 4220 with four or more genes). The focal analysis also recovered fewer COGs with high confidence (PP>0.9) at the root (see gray diagonal fields in table S13). Of the COGs recovered with high confidence at the root in the focal analysis, the majority (63%-78%) were also recovered with high confidence at the root in the secondary analysis (see gray off-diagonal fields in table S13). Considering COGs recovered with high confidence at the root, root 3 of the focal analysis is again the least congruent with the other gene sets, and at the same time is the one with the smallest number of COGs recovered at the root. Importantly, all of the ancestral pathways discussed in the main text were recovered at the root with moderate to high PP support in both datasets, as indicated (table S7).

Fig. S1: Maximum likelihood outgroup-rooted bacterial phylogeny. The maximum likelihood phylogeny obtained under the best-fitting LG+C60+R8+F model on a concatenation of 30 marker genes shared between Bacteria and Archaea. The bacterial root (marked by a black arrow) separates CPR, Cyanobacteria +Margulisbacteria, and Chloroflexota+Dormibacterota from the rest of the bacterial tree, but this position has poor bootstrap support and a range of alternative hypotheses could not be rejected statistically; note also that a basal position for DPANN within Archaea (1, 81) could not be rejected using an Approximately Unbiased (AU) test (table S2). FCB are the Fibrobacterota, Chlorobia, Bacteroidota and related lineages; PVC are the Planctomycetota, Verrucomicrobiota, Chlamydia and related lineages; DST are the Deinococcota, Synergistota, and Thermotogota; ACD are Aquificota, Campylobacterota, and Deferribacterota; FA are Firmicutes and Actinobacteriota. Branch supports are ultrafast bootstraps, as indicated by the colour key. Branch lengths are proportional to the expected number of substitutions per site.

Effects of extinctions

Fig. S2: Lineage extinction has a negligible impact on the accuracy of DTL estimates using ALE. We simulated species trees and gene families in the context of varying levels of lineage extinction, then evaluated the correlation between inference error (inferred - simulated numbers of events / family size). There is a statistically significant but quantitatively small impact of extinction on inference accuracy, with higher proportions of extinct lineages resulting in increased inference error (Pearson's r ; values and p -values indicated in figure). Inference errors remain centred on 0 even when most lineages are extinct.

Simulation of realistic families

Fig. S3: Simulating gene families with numbers of DTL events similar to the real data. To align our simulations with empirical reality, we simulated a large number of gene families and then subsampled families for analysis by matching them against a random subsample of the empirical families (see Supplementary Materials and Methods). The result was a simulated dataset with distributions of DTL events close to the real data. On top, the gain events (y axis) against the loss events for simulated and real families. On the bottom, the ratio of T/D against losses. The simulated families resemble the real families in terms of DTL rates and rate ratios.

Fig. S4: Accuracy of ALE inference of duplication (D), transfer (T) and loss (L) events. The plot shows the accuracy ((inferred number of events - real number of events) / family size) of ALE inferences on a set of 2000 simulated families with observed numbers of DTL events similar to the real data (see fig. S3). ALE accurately estimates the correct number of DTL events for simulated families, with mean errors of all three types of events close to zero (D $\sim$ -0.002, T $\sim$ -0.0210, L $\sim$ 0.048).

Inference of evolutionary events

Duplications

Transfers

Losses

Fig. S5: Inference errors are most common in small gene families. We ordered the families by number of genes and computed the rolling mean of the difference between the real (simulated) fraction of events and the fraction inferred by ALE. Most inference errors occur in small families.

Inference of evolutionary events

Fig. S6: Performance of ALEml_undated on simulated datasets. These scatterplots visualise the relationship between inferred and inferable numbers of duplications, transfers, losses and speciations on simulated gene families. Inferable events are events that leave traces in gene trees; for example, gene losses in extinct lineages cannot be inferred from gene trees. “Speciations” refer to the number of speciation events that a gene lineage passes through. The dotted lines correspond to perfect accuracy. ALE accurately recovers all types of events over a broad range of rates, although the inferred number of events is underestimated when the true number of events is high. This is most evident for losses, which may be because multiple recent losses can be misinterpreted as a single loss in the common ancestor of the affected branches.

Fig. S7: Comparison of errors in inferred DTL events from simulated data suggests that ALE can distinguish duplications, transfers and losses from their impact on gene tree topologies. We investigated the relationship between inference errors for duplications (D), transfers (T), losses (L) and speciations (S), with error calculated as the difference between the inferred and true (simulated) number of events per gene family. A negative correlation between errors (for example, over-estimation of the number of transfers coupled with under-estimation of duplications or losses) would indicate that ALE tends to mistake one kind of event for another. There is no evidence of negative correlation between any types of events (all correlations were weakly to moderately positive, indicating that co-underestimation is the most common type of error, as observed for the independent events in fig. S6 above), suggesting ALE can distinguish events based on the patterns they leave in gene trees. As expected, errors in speciations and losses are positively correlated: if a gene is inferred to have originated more recently in the species tree than is in fact the case (that is, if the number of speciations it has experienced are underestimated), then the number of losses required to explain the gene tree will also be underestimated.

Fig. S8: Impact of genome size heterogeneity (variation in DTL rates across the tree) on ALE inference accuracy. To investigate the effect of lineage-specific biases in DTL rates on ALE inference accuracy, we performed a simulation in which gene family originations were constrained to occur at the same nodes as in the real data. This procedure results in a set of simulated gene families that recapitulate the variation in genome size and gene content observed in the real data (for example, with a clade of small-genome CPR). We then compared inference accuracy (simulated - inferred events / family size) for D, T and L rates to a control simulation in which originations occurred at random, resulting in homogeneous genome sizes and gene contents across the tree. The impact of genome size heterogeneity is negligible, and all errors in both settings were centred on 0.

Fig. S9: Maximum likelihood unrooted bacterial phylogeny under the best-fitting substitution model (LG+C60+R8+F) following removal of the 20%-80% most compositionally heterogeneous sites. While the trees were similar overall, the position of Fusobacteriota was unstable, either branching as in the focal tree (40%, 60%, 80% of most compositionally heterogeneous site removed) or with (DST (20% of sites removed)). Sites were identified and removed using Alignment Pruner. (a) 20% most compositionally heterogeneous removed, with 14580/18234 sites remaining following site stripping; (b) 40% most compositionally heterogeneous removed, with 10941/18234 sites remaining following site stripping; (c) 60% most compositionally heterogeneous removed, with 7294/18234 sites remaining following site stripping; (d) 80% most compositionally heterogeneous removed, with 3647/18234 sites remaining following site stripping; Branch supports are ultrafast bootstraps, branch lengths are proportional to the expected number of substitutions per site.

Fig. S10. Duplication (D), transfer (T) and loss (L) rates for MCL families. a) scatter plots of T vs D, L vs T and L vs D rates against each other, with ratios above ten falling outside (below and to the left) of the continuous black lines. The plots show that extreme rate ratios typically result from very low rates. b) plotting per-family DTL rates versus the number of species with genes shows that, with the exception of the D rate, very small rates are only inferred for gene families with very limited species representation. This result implies that there exist some broadly-distributed species gene families that do not fix duplicates, whereas broadly-distributed families typically undergo at least some transfers and losses.

Fig. S11. Duplication (D), transfer (T) and loss (L) rates for COG families. a) Scatter plots of T vs D, L vs T and L vs D rates against each other, with ratios above ten falling outside (below and to the left) of the continuous black lines. The plots show that extreme rate ratios typically result from very low rates. b) Plotting per-family DTL rates versus the number of species with genes shows that, with the exception of the D rate, very small rates are only inferred for gene families with very limited species representation.

Figure S12. Impact of gene family filtering on ALE root likelihoods. We ranked gene families according to various criteria, then performed analyses to evaluate the impact on root likelihoods of excluding (filtering out) an increasing proportion of the top-ranked gene families for each criterion. The procedure represents a phylogenomic equivalent of site filtering in traditional phylogenetics, and is a useful way to explore the kinds of gene families that support each root position. Each analysis was performed on both the MCL and COG gene family datasets. For each analysis, the top panel plots ΔLL , the difference in log likelihood between a given root and the maximum likelihood root, as a function of filtered gene families. Positive values of ΔLL indicate that a given root has a better likelihood than the overall ML root at that percentile. The second panel shows the distribution of that test quantity, and the value of the quantity at each percentile of gene families filtered. Panel (a) maps the colours used in subsequent panels to the position of that root on a schematic unrooted tree; for example, the root between Gracilicutes and the rest of Bacteria corresponds to the light green line in panels (b)-(s). Removal of gene families by (b-c) species representation (families present in the most species filtered first); (d-e) species representation (families present in the fewest species filtered first); (f-g) highest copy number; (h-i) highest duplication (D) rate; (j-k) highest transfer (T) rate; (l-m) highest loss (L) rate; (n-o) highest rate ratio; (p-q) highest verticality; (r-s) lowest verticality; (t-u) highest transfer/loss (T/L) rate ratio.

Fig. S12(a). Mapping of root likelihoods to branches on the bacterial tree.

Fig. S12(b). Filtering MCL families ranked by representation in largest number of species

Fig. S12(c). Filtering COG families ranked by representation in largest number of species

Fig. S12(d). Filtering MCL families ranked by representation in smallest number of species

Fig. S12(e). Filtering COG families ranked by representation in smallest number of species

Fig. S12(f). Filtering MCL families ranked by highest mean copy number

Fig. S12(g). Filtering COG families ranked by highest mean copy number

Fig. S12(h). Filtering MCL families ranked by D rate, high to low

Fig. S12(i). Filtering COG families ranked by D rate, high to low

Fig. S12(j). Filtering MCL families ranked by T rate, high to low

Fig. S12(k). Filtering COG families ranked by T rate, high to low

Fig. S12(l). Filtering MCL families ranked by L rate, high to low

Fig. S12(m). Filtering COG families ranked by L rate, high to low

Fig. S12(n). Filtering MCL families ranked by largest DTL rate ratio

Fig. S12(o). Filtering COG families ranked by largest DTL rate ratio

Fig. S12(p). Filtering MCL families ranked by verticality, high to low

Fig. S12(q). Filtering COG families ranked by verticality, high to low

Fig. S12(r). Filtering MCL families ranked by verticality, low to high

Fig. S12(s). Filtering COG families ranked by verticality, low to high

Fig. S13: Two rooted topologies from the secondary, GTDB-independent analysis that could not be rejected by the AU test, from ALE analysis incorporating genome completeness. AU p-values are 0.973 for tree (a) and 0.064 for tree (b). Both trees are in agreement with each other and with the focal analysis in placing the root between Terrabacteria and Gracilicutes, but disagree in the placement of the "FASST" taxa comprising Fusobacteriota, Aquificota, Synergistota, Spirochaetota and Thermotogota. D-T stands for Deinococcota-Thermus; "Deltaproteobacteria" is Desulfuromonadota, Desulfobacterota, Bdellovibrionota, and Myxococcota.

Relative dating simulations

Fig. S14: Performance of relative dating approach on simulated data. We performed a relative dating analysis on a simulated dataset using the same number of gene families and rates as in the real dataset (11272), on a tree with 265 species. We applied the same pipeline used on the real dataset to obtain the relative dates. In the simulated dataset we obtain 16,943 highly supported constraints (found at least in 95 of the bootstrap replicates, see (36)), roughly twice as many as in the real dataset. The vertical blue lines represent the uncertainty in the ranking of the different speciation nodes. This approach correctly recovers the ranking of 98.4% of the nodes of the species tree, with only 4 being outside the distribution.

Fig. S15: The relationship between verticality and gene family size. Most gene families have experienced many transfers. Verticality varies with gene functional class, but families with very low transfer rates are small; these might represent young families that have not yet had enough time to experience gene transfer.

COG families

Genome size (mb)

Fig. S16: Evolution of COG family repertoires and inferred genome size over the bacterial tree. (a) The inferred number of COG family members and (b) inferred genome size at each internal node of the tree. Genome sizes were predicted from the relationship between COG family members and genome size among extant Bacteria (LOESS regression). Circle diameter is proportional to family number or genome size. FCB are the Fibrobacterota, Chlorobia, Bacteroidota, and related lineages; PVC are the Planctomycetota, Verrucomicrobiota, Chlamydiia, and related lineages; DST are the Deinococcota, Synergistota, and Thermotogota; ACD are Aquificota, Campylobacterota, and Deferribacterota; FA are Firmicutes and Actinobacteriota. The figure depicts inferences for root 1 (as shown in Fig. 1B); the data for all three roots are provided in GenomeSizeTable.tsv in the Online Data Supplement (80).

Fig. S17: Metabolic map of the central metabolic pathways inferred in the last bacterial common ancestor (LBCA) and a selection of subsequent nodes. The reconstruction is based on genes that could be mapped to a given node with PP >0.5. The presence of a gene within a pathway is indicated as shown in the key. Annotations and PP values for KOs in this figure can be found in table S8. Annotations and PP values for all KOs can be found in table S7.

Fig. S18: Distribution of COG families from key metabolic pathways inferred to the last bacterial common ancestor (LBCA). The occurrence of COG families in the taxa sampled in this study are represented as percentage presence across phylogenetic clusters (phylum) based on a presence/absence table. COG families inferred to the given nodes and the tree possible root positions (see Methods) are represented by corresponding PP values ($PP > 0.5$). TCA=Tricarboxylic Acid Cycle, PPP=Pentose phosphate pathway, ASR=Assimilatory sulfate reduction, DSR=Dissimilatory sulfate reduction, LPS=Lipopolysaccharide. Table S7 lists the PP values for all COG occurrences across roots, nodes, and tips and the metabolic genes featured in this plot can be found in table S8. A full heat map for all COGs, and additional heat maps by COGs category, can be found in (80), in heatmaps.zip.

Supplementary Tables

KO number	Gene name	Annotation	Used in outgroup tree?
K03046	rpoC	DNA-directed RNA polymerase subunit beta'	y
K03043	rpoB	DNA-directed RNA polymerase subunit beta	
K02337	dnaE	DNA polymerase III subunit alpha	y
K03070	secA	Protein translocase subunit SecA	
K01873	VARs, valS	Valine--tRNA ligase	
K02335	polA	DNA polymerase I	y
K01872	AARS, alaS	Alanine tRNA ligase	
K02469	gyrA	DNA gyrase subunit A	
K00962	pnp, PNPT1	Polyribonucleotide nucleotidyltransferase	y
K02355	fusA, GFM, EFG	Translation elongation factor G	y
K01972	E6.5.1.2, ligA, ligB	DNA ligase NAD	
K03702	uvrB	Excinuclease ABC subunit B	
K02470	gyrB	DNA gyrase subunit B	
K04077	groEL, HSPD1	Molecular chaperone GroEL	
K02313	dnaA	Chromosomal replication initiator protein	
K02314	dnaB	Replicative DNA helicase	
K02433	gatA, QRSL1	aspartyl-tRNA(Asn)/glutamyl-tRNA(Gln) amidotransferase subunit A	
K03076	secY	Protein translocase subunit SecY	y
K04485	radA, sms	DNA repair protein RadA/Sms	
K02112	ATPF1B, atpD	F-type H ⁺ /Na ⁺ -transporting ATPase subunit beta	y

K03590	ftsA	Cell division protein FtsA	
K02358	tuf, TUFM	Elongation factor Tu	y
K06942	ychF	Redox Regulated ATPase YchF	
K00927	PGK, pgk	Phosphoglycerate kinase	
K01889	FARSA, pheS	Phenylalanine--tRNA ligase alpha subunit	
K03551	ruvB	Holliday junction branch migration DNA helicase RuvB	y
K04485	radA, sms	DNA recombination repair protein RecA	y
K02835	prfA, MTRF1, MRF1	Peptide chain release factor 1	
K02886	RP-L2, MRPL2, rplB	50S ribosomal protein L2	y
K01803	TPI, tpiA	Triose phosphate isomerase	y
K03438	mraW, rsmH	16S rRNA (cytosine1402-N4)-methyltransferase	
K00554	trmD	tRNA (guanine37-N1)-methyltransferase	
K02863	RP-L1, MRPL1, rplA	50S ribosomal protein L1	y
K03685	rnc, DROSHA, RNT1	Ribonuclease III	
K02967	RP-S2, MRPS2, rpsB	30S ribosomal protein S2	y
K02982	RP-S3, rpsC	30S ribosomal protein S3	
K02906	RP-L3, MRPL3, rplC	50S ribosomal protein L3	y
K03470	rnhB	Ribonuclease HII	
K01358	clpP, CLPP	ATP dependent Clp protease proteolytic subunit	
K06187	recR	Recombination protein RecR	
K15034	yaeJ	Aminoacyl tRNA hydrolase, ribosome-associated protein	
K02931	RP-L5, MRPL5, rplE	50S ribosomal protein L5	y
K02933	RP-L6, MRPL6, rplF	50S ribosomal protein L6	y

K02601	nusG	Transcription termination antitermination protein NusG	
K02988	RP-S5, MRPS5, rpsE	30S ribosomal protein S5	y
K02992	RP-S7, MRPS7, rpsG	30S ribosomal protein S7	y
K03664	smpB	SsrA binding protein	
K02838	frr, MRRF, RRF	Ribosome recycling factor	
K02867	RP-L11, MRPL11, rplK	50S ribosomal protein L11	
K02878	RP-L16, MRPL16, rplP	50S ribosomal protein L16	y
K02871	RP-L13, MRPL13, rplM	50S ribosomal protein L13	y
K02994	RP-S8, rpsH	30S ribosomal protein S8	y
K02948	RP-S11, MRPS11, rpsK	30S ribosomal protein S11	y
K02952	RP-S13, rpsM	30S ribosomal protein S13	y
K02935	RP-L7, MRPL12, rplL	50S ribosomal protein L7/12	
K02996	RP-S9, MRPS9, rpsI	30S ribosomal protein S9	
K02874	RP-L14, MRPL14, rplN	50S ribosomal protein L14	y
K02887	RP-L20, MRPL20, rplT	50S ribosomal protein L20	
K02946	RP-S10, MRPS10, rpsJ	30S ribosomal protein S10	y
K02965	RP-S19, rpsS	30S ribosomal protein S19	y
K02956	RP-S15, MRPS15, rpsO	30S ribosomal protein S15	
K02518	infA	Translation initiation factor IF 1	

Table S1: 62 orthologous genes used to infer the species tree, with those used in the outgroup rooting analysis indicated.

Root hypothesis	log-likelihood difference to ML	p-value	Study
Observed outgroup root (Fig. S1)	0	0.58	This study (ML tree)
Between Firmicutes and Actinobacteriota	-3.3	0.50	(26)
Deinococcota, Synergistota and Thermotogota basal*	-3.7	0.52	
Planctomycetota basal	-4.9	0.49	(8)
Chloroflexota basal	-6.2	0.52	(9)
CPR basal	-16.7	0.37	(11, 16)
DPANN basal within archaeal outgroup	-19.4	0.37	(1, 20)
Fusobacteriota basal	-24.2	0.33	
Between Gracilicutes and Terrabacteria	-24.5	0.33	This study (ALE root, see below)

Table S2: Support for published hypotheses using outgroup rooting. *Our unrooted topology was incompatible with some published hypotheses, including a clade of Thermotogales and Aquificales at the root (6, 7).

Root	<i>p</i> -value	Study
CPR basal	2e-04	(11,16)
Chloroflexota (Chloroflexi) basal	1e-41	(9)
Between Firmicutes and Actinobacteriota	9e-05	(26)
Thermotogota/ Synergistota/ Deinococcota basal*	0.004	
Planctomycetota (Planctomycetes) basal	2e-26	(8)

Table S3: Support for published rooting hypotheses from our outgroup-free analyses. *Our unrooted topology was incompatible with some published hypotheses, including a clade of Thermotogota and Aquificota at the root (6, 7).

Table S4: AU p -values for all tested roots in ALE analysis (Excel-formatted spreadsheet)

Root name	LLs	AU
Fusobacteria root (398)	-7.3	0.589
Fusobacteriota on Terrabacteria side (527)	7.3	0.519
Fusobacteriota on Gracilicutes side (528)	13.6	0.432
Fusobacteriota and Spirochaetota on Terrabacteria side (520)	32.1	0.251
DST root (464)	103.7	0.008
Cyanobacteria on Gracilicutes side (517)	215.8	1e-09
Dormi/Chloroflexi+CPR (510)	372.2	7e-06
CPR root (496)	425.2	2e-67
Dormibacterota/Chloroflexota (505)	630.7	5e-05
Omnitrophota/Verrucomicrobiota/Planctomycetota (492)	674.4	2e-04
Fibrobacterota/Bacteroidota/Marinisomatota (511)	1000.7	3e-71

Table S5: AU-test results for an ALE root analysis on the focal dataset using 3595 COG families.

Root branch	1	2	3
Median singleton support per branch	999.09	998.47	999.21
Singleton support for branches subtending the root	98.259, 140.45	91.56, 151.09	95.25, 117.16
Mean verticality	0.641	0.641	0.642

Table S6: Singleton support (the number of genes that evolve vertically from one end of a branch to the other) on the credible set of rooted trees. Root numbers correspond to the three root branches depicted in Fig. 1(b).

Table S7: Protein family annotations (COG and KO) and root presence posterior probabilities (PPs) for key pathways used in ancestral reconstruction (Excel-formatted spreadsheet).

Table S8: COG families lost on the CPR stem (Excel-formatted spreadsheet).

Table S9: Species names and accessions of the genomes used in this study (Excel-formatted spreadsheet).

COG	Annotation	Number of families	Median transfer propensity	Mean transfer propensity
V	Defense mechanisms	32	0.4720	0.4361331747
T	transduction	88	0.4201	0.3955678194
G	Carbohydrate	186	0.4086	0.3984357156
Q	S e c o n d a r y metabolites	73	0.4036	0.3926403801
L	Replication	179	0.4012	0.3849087174
P	Inorganic ion	199	0.4011	0.3884050415
O	Post-translational modification	123	0.3986	0.3868587153
K	Transcription	131	0.3922	0.3659200418
I	Lipid	77	0.3875	0.3783924085
C	Energy	239	0.3854	0.3724628343
M	Cell wall/membrane	141	0.3845	0.3746415978
H	Coenzyme	165	0.3766	0.3704811363
E	Amino acid	226	0.3765	0.3684339442
F	Nucleotide	102	0.3577	0.3629977725
N	Cell motility	70	0.3174	0.3198036954
D	Cell cycle	45	0.3150	0.3094970449
U	Intracellular trafficking	83	0.3146	0.3119251964
J	Translation	177	0.3093	0.312946966

Table S10: Median and mean transfer propensity $T/(V+T)$ by COG functional category.

COG category	Number at root (flat prior)	Number at root (ML prior, PP >= 0.95)	Number at root (ML prior, PP >= 0.8)	Number of families in category	ML O_R
J	3	125	157	177	6739.434
F	1	57	89	102	5294.183
L	4	42	100	179	1468.033
H	3	26	85	165	1393.54
-	0	0	1	2	1290.933
N	1	18	38	70	1270.758
E	2	45	116	226	1146.195
D	2	7	16	45	850.0477
P	4	16	64	199	650.1995
M	1	15	49	141	625.6504
C	4	22	57	239	581.1914
G	3	11	43	186	513.1411
I	0	1	12	77	457.9165
U	1	3	12	83	364.7495
K	1	3	10	131	260.9843
O	0	1	9	123	256.3592
S	6	9	69	1374	163.0099
B	0	0	1	5	126.1288
T	0	1	3	88	125.0658
V	0	0	0	32	76.72498
Q	0	0	0	73	1.196214
A	0	0	0	4	1
Z	0	0	0	2	1

Table S11: Estimated root origination rates and root presences by COG functional category.

Clade	No. of taxa sampled
Firmicutes	25
Actinobacteriota+Cyanobacteria+Chloroflexota	35
CPR	125
FCB+PVC+Elusimicrobiota	35
Proteobacteria	50
“Deltaproteobacteria” +Nitrospirota+Acidobacterota+Aquificota	30
FASST+environmental lineages	25
New Phyla (Parks et 2018)	17

Table S12: Number of taxa sampled from each clade in the GTDB-independent analysis.

	Focal root 1	Focal root 2	Focal root 3	Secondary root 1	Secondary root 2
Focal root 1	475/3782	0.94	0.89	0.53	0.50
Focal root 2	403/3782	502/3782	0.9	0.53	0.51
Focal root 3	271/3782	265/3782	272/3782	0.49	0.48
Secondary root 1	364/3677	375/3677	210/3677	1383/4220	0.96
Secondary root 2	360/3677	378/3677	213/3677	1294/4220	1468/4220

Table S13. Comparison of ancestral gene contents between the focal and secondary datasets for different roots. Values with a white background correspond to correlation coefficients between PPs for COGs shared between two candidate roots (Pearson's correlation, all values highly significant with $p < 10^{-16}$); values with gray background correspond to the number of shared COGs with PP>0.9 / total number of shared COGs, with the diagonal fields giving corresponding values for each root on its own.